The Congested Assignment Problem*

Anna Bogomolnaia† and Hervé Moulin‡

May 2023

Abstract

We must assign n agents to m posts subject to negative congestion. If congestion is anonymous (each agent adds one unit) a canonical test of ex ante fairness guarantees to everyone one of their top n out of $n \times m$ feasible allocations. We call an assignment competitive if, taking the congestion at each post as given, everyone is assigned to one of their best posts. If it exists the competitive assignment is unique, efficient and ex ante fair.

Randomising the assignment under cardinal vNM utilities ensures that every problem has a unique competitive congestion profile, implemented by a mixture of deterministic assignments rounding up or down the competitive congestion, and approximately ex ante fair, efficient and welfare equivalent.

If congestion is weighted (agent-specific) ex ante fairness still guarantees an allocation in the top (1/m) quantile of feasible allocations, and under randomisation competitiveness is still an operational solution concept.

Keywords: assignment, congestion, guarantee, competitive, envy free, free mobility game

JEL classification: D47, D63

**Acknowledgments*: We thank Felix Fischer, Ioannis Caragiannis, Igal Milchtaich and Sinan Ertemel for critical comments and references.

†Adam Smith Business School, University of Glasgow and Centre d'Economie de la Sorbonne, CNRS, Paris

‡Adam Smith Business School, University of Glasgow

1 introduction

Congestion affects the allocation of any commodity consumed under partially rivalry: traffic on roads and the internet, public utilities and private services, etc.. For just over a century it is a staple of microeconomic analysis. Pigou's work on the taxation of externalities([30]) inspired much research looking for the welfare-optimal taxation of congestion, for instance peak load pricing of utilities ([31], [8]),Vickrey's optimal toll road ([38]) and the far reaching Vickrey-Clarke-Groves pricing mechanisms ([17]).

A different operations research and game theory literature starting with Wardrop's transportation model ([13]) analyses congestion in the decentralised *free mobility* regime: each agent chooses her service (what route to use, which store to visit, etc..) and the resulting non cooperative equilibrium determines the congestion. Such games include the seminal instance of potential games ([32], [25]) and inspire the definition of the price of anarchy[1] ([23], [33])

We discuss here an assignment model that appeared first in the literature on congestion games ([24]): we are given a set of heterogenous items subject to congestion that we call "posts"; each agent must be assigned to a post, and each post can take any amount of congestion, from remaining empty to hosting everyone. But instead of allowing agents to choose their post in non cooperative fashion, we take a normative viewpoint: what could be a fair and efficient assignment? So in spirit we are closer to the microeconomic search for optimal taxation but with a key difference: cash transfers are ruled out, agents cannot be compensated to accept an inferior post, or pay to be assigned at a popular one.

Examples of moneyless assignment problems where congestion is an important consideration include the allocation of jobs to busy heterogenous servers, of workers to shared office spaces, patients to hospitals, students to crowded classes, students to crowded schools, or messages to routes in a centralised communication network. In those examples the free mobility "choose your own post" approach is typically impractical and possibly inefficient. It is also unpalatable because it gives an unfair advantage to agents who happen to play the game better: see the discussion of the

[1]Measuring the welfare loss in the worst Nash equilibrium of the game; see e. g. **[16],** [34].

perverse strategic implications of the Boston school assignment mechanism ([14]).

Matching and assignment models involving "many agents to one position" routinely incorporate *hard* congestion constraints in the form of upper or lower bounds on the filling of each post: maximal capacity of a class ([9]) or a school ([1]), minimal or maximal quotas for some subsets of students ([3], **[18], [26], [20])** etc.. But when these constraints do not bite or are altogether absent, *soft* congestion still impacts agents' welfare and choices. Parents will accept more crowded classes if the school's academic context is better, or vice versa: the recent paper [29] on congested school choice (on which more in the next section) offers clear evidence of this point.

Our definition of fairness, in line with seven decades of microeconomic literature on the division of privatcommodities (e. g. [27]), is two-fold. First, ex ante fairness focuses on the worst case welfare level that can be guaranteed to each agent, based on their own trade-offs between posts and congestion but no other information about the preferences of the other potentially adversarial participants. Second, ex post fairness adapts the familiar envy-freeness property ([15], [37]), achieved by randomisation like in the iconic assignment model (aka unilateral matching) to which we are adding the congestion dimension.

Our punchlines: the ex ante test "top-fairness" relies only on ordinal preferences and severely reduces, sometimes to a singleton, the set of acceptable assignments. If the agents have cardinal vNM utilities the ex post "competitiveness" test singles out an essentially unique efficient assignment, also top-fair and efficient: it is a fair and efficient way to manage congestion.

Notwithstanding the fact that it requires elicitation of vNM cardinal utilities rather than simpler ordinal preferences, we submit that our competitive analysis applies to a rich family of congested assignment problems and delivers an efficient solution built upon compelling fairness principles.

1.1 overview of the results

We review the main modeling choices and findings of the paper. In sections 2,3 congestion is anonymous (aka unweighted): each agent adds one unit of congestion at any post. Congestion is weighted in section 4: different agents may bring a different amount of congestion, but the same to each post.

In the anonymous case section 2.1 defines the model with ordinal preferences; in particular an agent's allocation is (a, s_a): to share post a with $s_a - 1$ other agents. We have m posts and n agents hence $m \times n$ feasible allocations; the only restriction on preferences over the latter is that they decrease strictly in s_a.

The point of ex ante fairness is to make the wost case welfare of each agent as high as possible.[2] In the *standard* assignment problem (one-to-one mapping of n agents to n posts) if all agents have identical strict preferences, one of them must be assigned to the unanimously worst post: therefore the worst case welfare is the lowest of the n possible ones. When the number m of posts is not related n and posts can be congested, the worst case is still the n-th best welfare, but this time out of $m \times n$ possible levels. We call *top-n-fair* an assignment respecting this constraint for every agent: their existence is guaranteed by Lemma 1 section 2.2.

The second fairness property, *competitiveness*, refines the familiar envy-freeness (Definition 2 section 3.1): if assigned to post a with congestion s_a I weakly prefer (a, s_a) to each allocation (x, s_x) realised by the assignment under consideration, including to $(y, 1)$ if the post y is empty ($s_y = 0$). Taking the congestion price of post x as the largest of s_s and 1, the aggregate competitive demand confirms the proposed assignment (or contains it if some demands are multivalued).

If it exists in the ordinal model, the competitive assignment is compelling: fair in the top-n-fair and competitive sense, efficient and unique both in terms of congestion and welfare (Proposition 1). But problems where none exists abound, so we use in section 3.2 a familiar randomisation technique to produce an operational concept of solution.

There are two ways to randomise the standard assignment model to ensure that an envy-free and efficient allocation always exists. Most popular in this century is the ordinal way: we extend each deterministic ordinal preference ordering to its (incomplete) stochastic dominance ordering of random allocations; the resulting set of envy-free and efficient assignments is typically large (in the sense of having a non

[2] Recall Steinhaus' cake-cutting mechanism ([36]): a truthtelling agent secures a share worth at least $\frac{1}{n}$-th of their valuation for the entire cake, irrespective of the reports of the other $n-1$ participants; this is as much as we can guarantee to everyone.

empty interior) and the Probabilistic Serial (PS) algorithm delivers a simple selection ([7], [21], [4]). In the congested model this ordinal definition of competitiveness fails because it is vastly underspecified: if we fix an arbitrary congestion profile (s_x) we can force the PS algorithm to implement it, but it is not clear what deeper principle should guide the choice of that profile. See section 3.2.2.

We adopt instead the second approach to randomisation, proposed in the standard model more than forty years ago in ([19]). Assume the agents have von Neuman Morgenstern (vNM) cardinal utilities over allocations and construct a competitive equilibrium where each agent uses an equal budget of fiat money to buy probabilistic shares of the posts at the given price; this equilibrium is not always unique though one expects a finite number of isolated assignments.

We introduce the cardinal vNM utilities in section 3.2: a randomized assignment generates a *fractional* (real valued) expected congestion σ_x at each post x and an agent's competitive demand is simply the set of her preferred allocations (x, σ_x) (there is no fiat money to distribute among posts). In this now convex range of congestion profiles the existence of a congestion profile (σ_x) confirmed by the aggregate competitive demand follows from a standard fixed point argument. Remarkably there is a unique competitive profile (σ_x^c) (Lemma 2 section 3.2). It is implemented by a lottery over finitely many deterministic assignments: in each of these congestion rounds the competitive congestion up or down to the nearest integer; they (the deterministic assignments) are approximately top-n-fair, competitive and efficient; and have appproximately identical utilities (Lemma 3 and Theorem 1 section 3.2). Section 3.2.1 illustrates our main result in some examples.

Many of the above results apply to the more general and technically more demanding model in section 4 where each agent i brings the amount w_i of congestion to her post. In the ordinal model the definition of ex ante fairness is more subtle because my guaranteed welfare level only depends upon my own preferences and the total congestion W in the problem, not on how W is divided or even on the total number n of agents (Definition 5 section 4.1). It still guarantees that everyone ends up in the top $\frac{1}{m}$ quantile of his preferences: Proposition 2 (section 4.1) requires more work than Lemma 1 in the anonymous case.

With deterministic assignments and ordinal preferences the competitive analysis mostly mimicks the anonymous case: Proposition 3 section 4.2 only adds some caveats to Proposition 1. Introducing vNM cardinal utilities and randomisation in section 4.3, it still identifies a unique competitive congestion profile provided we rule out single occupancy posts (Lemma 4 section 4.3). We can then implement this congestion by a lottery over deterministic assignments in which each agents is always assigned to a post in his competitive demand (Lemma 5). But unlike in the anonymous congestion case these assignments no longer approximate the competitive congestion because the difference between individual weights is unbounded.

Section 5 repeat the take home points and state two open questions. Section 6 gives some elaborate examples omitted from the main body of the paper.

1.2 more related literature

1) Two independent recent papers introduce congestion in school choice. Closest to ours is Phan et al. [29] adding the congestion dimension to the school choice model ([1]): students' preferences over shools depend also on their crowding level measured by the per capita resources at each school. This is formally isomorphic to our own model. To manage crowding the paper adopts like us a market clearing definition of ex post fairness adapting envy-freeness like we do. But the school choice viewpoint adds hard capacity constraints as well as priorities over students for the schools: this qualifies envy-freeness as "no justified envy" and their novel concept of Rationing Crowding Equilibrium (RCE) is more complex than our competitive assignment; also the proof that a RCE exists is significantly more involved. Remarkably the RCEs have the semi-lattice structure from which emerges a maximal RCE student optimal and unique welfarewise, as in the standard school choice model.

Phan et al's model is deterministic and relies on ordinal preferences, so it cannot take advantage of the convexification offered by randomisation and vNM cardinal utilities. Still, up to a rounding argument (reminiscent of the one in our Lemma 3) in the definition of RCEs, Theorem 2 showing that the RCEs have the same crowding profile is analog to statement i) in our Proposition 1 and Lemma 2.

Copland [12] defines a school choice model, also deterministic and with ordinal preferences, where students have strict preferences over allocations (a, s); instead of envy-freeness he imposes what we call the free mobility equilibrium property to balance fairly the congestion across schools. This property implies our top-n-fairness (section 2.2.1) but is substantially more permissive than envy-freeness (section 3.1.1). Copland then proposes a variant of the Deferred Acceptance mechanism (DA with Voluntary Withdrawals) to take into account the schools' prioritiies and capacity comstraimts. So the formal similarities with our results are fewer than in [29] but the general viewpoint on congestion is still the same.

2) The free mobility (FM) games in our model are not potential games but they stand out for their powerful existence results of the Nash and strong equilibria ([24], [22]). Similar games play a key role in the hedonic model of coalition formation ([2], [5]) and local public goods (e. g. [6]).

Under anonymous congestion every problem in our model has at least one FM equilibrium; as mentioned one paragraph above, each FM equilibrium outcome is a top-n-fair assignment and each competitive assignment is a FM equilibrium outcome (the converse statements do not hold).

The connection is more tenuous if congestion is weighted because the FM game may not have any Nash equilibrium ([24]). If it does, the equilibrium outcome is still top-$\frac{1}{m}$-fair (Remark 3 concluding section 4.1).

3) The classic combinatorial optimisation problems known as bin packing or knapsack ([10]) discuss like us how to fill bins (posts) with indivisible balls (agents); but the balls there are just objects and the concern is about the "welfare" of the bins (e.g., to respect a capacity constraint, or to minimise its load) while we focus on the welfare of the balls and treat the bins as objects.

One exception is [11] where each ball has its own maximal acceptable congestion in each bin: this is exactly like our description of top-n- guarantees (section 2.2), with the difference that the caps are exogenous in [11] so that we may not be able to assign all balls to some bin. The paper shows the complexity of computing the maximal number of balls we can assign and evaluates the price of anarchy of the corresponding free mobility equilibrium.

2 the canonical guarantee

2.1 the ordinal model

We have m posts denoted $a, b, \cdots$, and their set is A. Each agent i in the finite set N of cardinality n must be assigned to some post a in A, which creates one unit of congestion at a.

An *assignment* of agents to posts is a (quasi) partition $P = (S_a)_{a\in A}$ of N where S_a is the set of agents assigned to post a; the sets S_a, S_b are mutually disjoint and at most $m-1$ of them can be empty. The corresponding congestion at post a is $s_a = |S_a|$ (the cardinality of S_a); we call $s = (s_a)_{a\in A}$ the *congestion profile* of P. We write the set of assignments as $\mathcal{P}(A, N)$, or simply $\mathcal{P}$.

We use the notation $[q]$ for the interval $\{1, \cdots, q\}$ in $\mathbb{N}$. Agent i's (transitive and complete) preference relation $\succeq_i$ bears on the set $A \times [n]$ of feasible allocations (a, s_a): i's assigned post a and the congestion s_a at that post. preferences are strictly decreasing in s_a and otherwise arbitrary. A problem $(A, N, \succeq)$ is defined by the choice of a preference profile $\succeq = (\succeq_i)_{i\in N}$.

Given agent i with preferences $\succeq_i$ we can order the $n \times m$ allocations z in $A \times [n]$ as a sequence $\{z^{*1}, z^{*2}, \cdots, z^{*n\times m}\}$ in such a way that $z^{*k} \succeq_i z^{*(k+1)}$ for all k, $1 \leq k \leq n\times m-1$. There is only one such sequence if the preferences $\succeq_i$ are strict, possibly more because of indifferences. A *n-prefix* of $\succeq_i$ is a subset of n allocations containing the first n allocations in one of the sequences just defined. Because welfare decreases strictly in congestion, if $z = (a, s)$ is in a certain prefix so does (a, s') for all $s' \leq s-1$: therefore a n-prefix of $\succeq_i$ can be represented by a vector $\lambda_i = (\lambda_{ia})_{a\in A}$ such that

$$\lambda_{ia} \in \mathbb{N} \cup \{0\} \text{ for all } a, \text{ and} \sum_{a\in A} \lambda_{ia} = n \tag{1}$$

It is the union of the sets $\{a\}\times[\lambda_{ia}]$ over all a in A, with the convention $[a]\times[0] = \varnothing$.

We write $\Delta^{\mathbb{N}}(A; n)$ for the set of n-prefixes defined by (1). Clearly when indifferences of the relation $\succeq_i$ allow several n-prefixes λ_i for $\succeq_i$, any two of these differ by at most 1 in each coordinate λ_{ia}.

For instance the n-prefix $\lambda_{ia} = n, \lambda_{ib} = 0$ for $b \neq a$ means that i (weakly) prefers post a to any other post x, no matter how congested x and a are . Next if $\lambda_{ia} = \lfloor \frac{n}{m} \rfloor$ or $\lceil \frac{n}{m} \rceil$ for all a (where $\lfloor x \rfloor$ and $\lceil x \rceil$ are the smallest and largest integer bounded above and below by x) this agent's priority is to minimise the congestion level at their assigned post (up to the necessary rounding up or down).

2.2 the top-n-guarantee

Ex ante fairness is the guarantee that each agent's allocation is in one of their n-prefixes.

Definition 1 *Given a problem $(A, N, \succsim)$ and a profile $\lambda = (\lambda_i)_{i \in N}$ of n-prefixes λ_i in $\Delta^{\mathbb{N}}(A; n)$, the assignment $P \in \mathcal{P}$ is top-n-fair iff*

$$s_a \leq \lambda_{ia} \text{ for all } a \in A \text{ and } i \in S_a \tag{2}$$

We write $\mathcal{P}(\lambda)$ for the set of such assignments and $\mathcal{C}(\lambda)$ for the set of congestion profiles s when P varies in $\mathcal{P}(\lambda)$.

Lemma 1 *In the problem $(A, N, \succsim)$ there exists at least one top-n-fair assignment P for any profile λ of n-prefixes.*

Proof by a simple greedy algorithm and an induction argument on n.

If for some post a we have $\lambda_{ia} = 0$ for all i we set $S_a = \varnothing$ and it remains to prove the claim on the residual problem $(A \diagdown \{a\}, N, \lambda)$. We clean up in this way all posts that nobody accepts so we assume from now on that $\max_i \lambda_{ia} \geq 1$ for all a.

Clearly if the Lemma holds for a given n, the existence of a top-n-fair assignment holds as well for all profiles $\widetilde{\lambda}$ weakly larger than $\lambda \in \Delta^{\mathbb{N}}(A; n)$ in al coordinates.

Pick any post a and order the caps $\lambda_{ia}, i \in N$ as $\lambda^{*1} \geq \lambda^{*2} \geq \cdots$. Write $\widetilde{k}$ for the largest k s.t. $\lambda^{*k} \geq k$ (well defined because $\lambda^{*1} \geq 1$) and pick a $\widetilde{k}$-prefix S_a of the ordering $i \succsim^a j \Leftrightarrow \lambda_{ia} \geq \lambda_{ja}$ of N. Then $\{\lambda_{ia}\}_{i \in S_a} = \{\lambda^{*k}\}_{1 \leq k \leq \widetilde{k}}$ and $\lambda_{ja} \leq \widetilde{k}$ for each $j \in N \diagdown S_a$ by definition of $\widetilde{k}$.

Assigning S_a to a meets inequalities (2) for a, and in the residual problem $(A \diagdown \{a\}, N \diagdown S_a, \lambda)$ we have $\sum_{b \in A \diagdown \{a\}} \lambda_{jb} \geq n - \widetilde{k} = |N \diagdown S_a|$ for all j. The induction assumption on n concludes the proof. ■

If all preferences $\succsim_i$ coincide, then in any assignment P and for any common

n-prefix λ_0, at least one agent gets at P a least preferred allocation in λ_0: this implies that the top-n-guarantee cannot be improved.

Remark 1 If $n = m$ and everyone cares first to be alone in their post then the common n-prefix is $\lambda_{ia} = 1$ for all a, and the top-n-fair assignments correspond to a standard one-to-one assignment problem.

We illustrate the power of top-n-fairness (2) in three examples. In the first two $\mathcal{C}(\lambda)$ is a singleton and we show the simple way to recognise this fact.

Example 1: $m = 2 \Longrightarrow |\mathcal{C}(\lambda)| = 1$. Set $A = \{a, b\}$ and pick a profile $\lambda = (\lambda_i)_{i \in N}$ of n-prefixes. Label the agents from 1 to n so that λ_{ia} decreases (weakly) with i, while λ_{ib} increases (weakly). Keeping in mind $\lambda_{ia} + \lambda_{ib} = n$, the integer $\widetilde{k}$ in the proof of Lemma 1 is defined by the inequalities $\lambda_{\widetilde{k}a} \geq \widetilde{k} \geq \lambda_{(\widetilde{k}+1)a} \Longleftrightarrow \lambda_{\widetilde{k}b} \leq n - \widetilde{k} \leq \lambda_{(\widetilde{k}+1)b}$.

We see that the maximal congestion at a compatible with top-n-fairness is $\widetilde{k}$ while at b it is $n - \widetilde{k}$: therefore $(\widetilde{k}, n - \widetilde{k})$ is the only top-n-fair congestion. Moreover if $\lambda_{\widetilde{k}} \neq \lambda_{(\widetilde{k}+1)}$ there is also a single top-n-fair assignment.

Example 2: $m = 4, n = 18$, *and* $|\mathcal{C}(\lambda)| = 1$. The eighteen agents are split in five types labeled α to ε, with from two to five agents in each type. Agents of a given type have the same 18-prefix λ_i but not necessarily identical preferences. The profile λ of 18-prefixes and the unique top-18-fair assignment P are as follows

prefixes

	a	b	c	d
$\alpha\alpha\alpha$	3	2	10	3
$\beta\beta\beta\beta$	7	1	7	3
$\gamma\gamma$	3	8	3	4
$\delta\delta\delta\delta\delta$	4	2	9	3
$\varepsilon\varepsilon\varepsilon\varepsilon$	4	0	7	7

top-fair assignmt P

a	b	c	d
$\beta\beta\beta\beta$	$\gamma\gamma$	$\alpha\alpha\alpha\delta\delta\delta\delta\delta$	$\varepsilon\varepsilon\varepsilon\varepsilon$

Check that we can fit top-18-fairly at most 4 agents at post a; at most 2 at post b; at most 8 at post c; at most 4 at post d. As $4 + 2 + 8 + 4 = 18$ we see that $\mathcal{C}(\lambda)$ contains at most the profile $(4, 2, 8, 4)$; but $\mathcal{C}(\lambda)$ is non empty by Lemma 1, so this profile is its sole element. In this example assigning all α-s and δ-s to c is the only way to fit 8 agents there and in turn $\mathcal{P}(\lambda) = \{P\}$.

More generally write $c^{mx}(a;\lambda)$ for the maximal number of agents we can fit λ-fairly at post a. Lemma 1 implies $\sum_{a\in A} c^{mx}(a;\lambda) \geq n$ for any λ. If $\sum_{a\in A} c^{mx}(a;\lambda) = n$ then all top-n-fair assignments have the same congestion profile $s_a = c^{mx}(a;\lambda)$ for all a. The converse property holds as well (we omit the easy proof):

$$|\mathcal{C}(\lambda)| = 1 \Longleftrightarrow \sum_{a\in A} c^{mx}(a;\lambda) = n \tag{3}$$

Examples where $|\mathcal{C}(\lambda)| = 1$ but $|\mathcal{P}(\lambda)| > 1$ abound, e. g. in Example 1 due to indifferences, or problem (11) section 3.1.

Example 3: $m = 3, n = 12$, *and* $|\mathcal{C}(\lambda)| = 15$. There are six agents in each of the two types; the following 12-prefixes allow three top-12-fair congestion profiles respecting the symmetries of the problem

$$\text{prefixes} \quad \begin{array}{cccc} & a & b & c \\ \alpha\alpha\alpha\alpha\alpha\alpha & 6 & 4 & 2 \\ \beta\beta\beta\beta\beta\beta & 2 & 4 & 6 \end{array} \quad \text{top-fair assignmts} \quad \begin{array}{ccc} a & b & c \\ \alpha\alpha\alpha\alpha\alpha\alpha & 0 & \beta\beta\beta\beta\beta\beta \\ \alpha\alpha\alpha\alpha\alpha & \alpha\beta & \beta\beta\beta\beta\beta \\ \alpha\alpha\alpha\alpha & \alpha\alpha\beta\beta & \beta\beta\beta\beta \end{array} \tag{4}$$

There are twelve more congestion profiles $s = (s_a, s_b, s_c)$ in $\mathcal{C}(\lambda)$: $(6,1,5)$; $(6,2,4)$; $(6,3,3)$; $(6,4,2)$; $(5,4,3)$; $(5,3,4)$, and six more by the symmetry exchanging α and β.

Counting the top-n-fair assignments: the first symmetric congestion in (4) yields a single assignment, but the second one already 36 by permuting agents within their types. Taking all profiles of $\mathcal{C}(\lambda)$ into account we find $|\mathcal{P}(\lambda)| = 229$.

Many of the top-12-fair assignments (but not all[3]) can be inefficient (Pareto dominated) as well. E. g., if in the lowest assignment of (4) four α-agents prefer $(a,4)$ to $(b,4)$ and the last two have the opposite preference then only one of the fifteen corresponding assignments of the α-s between posts a and b is efficient.

2.2.1 free mobility equilibria and top-n-fairness

The canonical decentralised management of congestion is the *free mobility* (FM) game: given a problem $(A, N, \succeq)$ each agent chooses a post a in A then consumes the allocation (a, s_a) where $s_a - 1$ is the number of other agents who chose a.

[3] If $P \in \mathcal{P}(\lambda)$ is efficient within $\mathcal{P}(\lambda)$, it is clearly efficient within the entire feasible set.

Every Nash equilibrium of the FM game is a top-n-fair assignment. To check this write $x = (x_i)_{i \in N} \in A^N$ a profile of strategies and $s(a|x)$ for the resulting congestion at post a. The equilibrium property for agent i is $(x_i, s(x_i|x)) \succeq_i (a, s(a|x) + 1)$ for all $a, a \neq x_i$: at such post a she *may* strictly prefer some or all in $\{a\} \times [s(a|x)]$ to her equilibrium allocation; adding to those the allocations in $\{x_i\} \times [s(x_i|x) - 1]$ that she definitely prefers, we see that at most $(n-1)$ other allocations can beat $(x_i, s(x_i|x))$, so the latter belongs to at least one n-prefix λ_i. Choosing one λ_i for each i, we conclude that x implements a top-n-fair assignment.

The main result in [24] is that for any problem $(A, N, \succeq)$ the FM game has at least one Nash equilibrium. It is easy to give examples where the set of FM equilibrium assignments is a strict subset of $\mathcal{P}(\lambda)$. In Example 3 the first assignment in (**??**) is a FM assignment only if each α (resp. β) prefers weakly $(a, 6)$ to $(b, 1)$ (resp. $(c, 6)$ to $(b, 1)$). It is equally easy to show that FM equilibrium assignments can be inefficient: see Example 5 in section 3.1.1.

Remark 2 There are two strategic interpretations of top-n-fairness in the direct revelation mechanism where agents report independently their preferences and the mechanism implements a reportedly top-n-fair assignment. First, every Nash equilibrium of such mechanism must be top-n-fair because the truthful report ensures this. Suppose next that agent i, clueless about other agents' preferences, wants to maximise his worst case welfare: reporting a truthful n-prefix of $\succeq_i$ *ensures this and no other report does. Indeed in the adversarial case where everyone else reports the same* λ_i *an assignment rule treating equals as equally as possible can give agent i anyone of the allocations* (a, λ_{ia})*, in particular a worst one in this set. Of course, any information about other participants' preferences can be used to advantage by our agent but at some risk.*

3 competitiveness

3.1 deterministic

We use the notation $z \vee y = \max\{z, y\}$.

Definition 2 *Fix a problem* $(A, N, \succeq)$. *The assignment P is competitive (Comp)* *iff*

$$\text{for all } a \in A \text{ and } i \in S_a : (a, s_a) \succeq_i (x, s_x \vee 1) \text{ for all } x \in A \tag{5}$$

If all posts are occupied ($s_x \geq 1$ for all x) property (5) is just envy-freeness: switching my allocation to yours is a virtual move that does not affect the congestion at your post (unlike an actual move in the FM game). Interpreting the congestion s_x at post x as the "price" of consuming x, my competitive demand of the allocation (x, s_x) does not affect its price.

Property (5) assigns a price of 1 to an empty post x and insists that our agent do not want to move there either: this is important because an assignment where all agents share a single post a is automatically envy-free, but can be the absolute worst assignment for everyone!

Proposition 1 *Fix a problem* $(A, N, \succeq)$.
i) All competitive assignments have the same congestion profile (except possibly at some posts occupied by at most one agent), and the same welfare profile.
ii) A competitive assignment is weakly efficient, and efficient if preferences are strict and/or if all posts are occupied.
iii) A competitive assignment is top-n-fair.

Proof At the congestion profile $s \in \Delta^{\mathbb{N}}(A; n)$ (see(1)) we define agent i's *competitive demand* as

$$D(i; s) = \{a | (a, s_a \vee 1) \succeq_i (x, s_x \vee 1) \text{ for all } x \in A\} \tag{6}$$

The assignment P is competitive if and only if $a \in D(i; s)$ whenever $i \in S_a$.
i) Unique congestion. We fix two different congestion profiles s, s^* coming from different competitive assignments $P = (S_x)_{x \in A}$ and $P^* = (S^*_x)_{x \in A}$.

Define the set $A^* = \{a \in A : s^*_a \vee 1 > s_a \vee 1\}$ and assume that A^* is non empty, which will lead to a contradiction. Note that in A^* we have $s^*_a > s_a, 1$.

Fixing an agent i we claim that if $D(i; s^*)$ intersects A^* at a then $D(i; s)$ must be a subset of A^*. If the claim fails there is some $b \in D(i; s)$ outside A^* such that for all a in A^*: $(b, s_b \vee 1) \succeq_i (a, s_a \vee 1) \succ_i (a, s^*_a)$ (strict preference because $s^*_a > s_a, 1$). By the choice of b we also have $s_b \vee 1 \geq s^*_b \vee 1$; therefore $(b, s^*_b \vee 1) \succeq_i (b, s_b \vee 1) \succ_i (a, s^*_a)$, and as a was arbitrary in A^* it follows that $D(i; s^*)$ cannot intersect A^*, contradiction.

Now for each $i \in \cup_{a \in A^*} S_a^*$ the claim says $D(i;s) \subseteq A^*$ therefore i is assigned to A^* by P. This implies $\sum_{a \in A^*} s_a \geq \sum_{a \in A^*} s_a^*$, contradicting $s_a^* > s_a$ in A^*. We conclude that A^* is empty.

So P, P^* must be such that $s_a^* \vee 1 = s_a \vee 1$ for all a: therefore $s_a^* \neq s_a$ can only happen when $\{s_a^*, s_a\} = \{0, 1\}$, as claimed in statement i). Moreover the competitive demands at s and s^* coincide.

A simple example with multiple competitive congestion profiles has all agents except 1 and 2 refusing the three posts a, b, c, while 1 and 2 refuse all but a, b, c and they are indifferent between $(a, 1), (b, 1)$ and $(c, 1)$: combining a competitive sub-assignment of $N \diagdown \{1, 2\}$ to $A \diagdown \{a, b, c\}$ with any assignment where agents 1,2 occupy two of a, b, c is competitive in the full problem.

Unique welfare. We fix an agent i in S_a and S_b^* and show that i is indifferent between the two assignments. From $s_a, s_b^* \geq 1$ and Comp we have

$$(a, s_a) \succeq_i (b, s_b \vee 1) \text{ and } (b, s_b^*) \succeq_i (a, s_a^* \vee 1) \tag{7}$$

If $s_a = s_a^*$ and $s_b = s_b^*$ we are done. If $s_a \neq s_a^*$ and $s_b = s_b^*$ statement i) implies $s_a = 1 > 0 = s_a^*$; then $s_b^* \geq 1$and (7) gives $(a, 1) \succeq_i (b, s_b) = (b, s_b^*) \succeq_i (a, 1)$ as desired. The last subcase $s_a \neq s_a^*$ and $s_b \neq s_b^*$ is just as easy.

ii) *Efficiency.* Assume, to the contrary that $P = (S_x)$ is competitive and Pareto inferior to $Q = (T_x)$. Say i assigned to a at P is assigned to b at Q (a, b are not necessarily distinct) and suppose that post b is occupied at P: $s_b \geq 1$. Then by Comp and the weak Pareto improvement we have

$$(b, s_b) \preceq_i (a, s_a) \preceq_i (b, t_b) \Longrightarrow s_b \geq t_b \tag{8}$$

and $s_b > t_b$ if agent i improves strictly at Q.

If all posts are occupied at P (8) implies $s = t$ and we have a contradiction. If instead some agent i goes from a at P to c at Q and c is empty at P, $s_c = 0$, we have

$$(c, 1) \preceq_i (a, s_a) \preceq_i (c, t_c) \Longrightarrow t_c = 1 \text{ and } (a, s_a) \simeq_i (c, 1) \tag{9}$$

This is a contradiction if preferences are strict, and we conclude again that P is Pareto optimal. Even if they are not we see that i is a weak Pareto optimum (not all agent benefit strictly).

The argument above explains how a competitive assignment can be Pareto inferior: in the situation (9) moving i from a to c *and changing nothing else* is a Pareto improvement to a new assignment $\widetilde{P}$ where agents in $S_a \diagdown \{a\}$ benefit strictly while all others are unaffected.[4]

Note that $\widetilde{P}$ is not competitive if $s_a \geq 2$ because i prefers strictly $(a, s_a - 1)$ to (a, s_a) and therefore to $(c, 1)$; but if $s_a = 1$ then $\widetilde{P}$ is competitive and welfare-wise indifferent to P.

As announced in statement *ii*) Pareto improving a competitive assignment can only happen if $\succeq_i$ is not strict and some post is unoccupied.

iii) *top-n-fairness.* A FM equilibrium is top-n-fair (section 2.2.1). Now a competitive assignment is clearly a FM equilibrium. ■

Proposition 1 speaks loudly in favour of a competitive assignment as an essentially single-valued, fair and efficient solution to our congested assignment problem. To illustrate its successful application and its limits, we start with Remark 1 just before Example 1. A standard *non congested* assignment problem obtains in our model when $n = m$ and every n-prefix is $\lambda_{ia} = 1$ for all a: then a competitive assignment exists only in the rare case where each agent has a different top post. In the second example with just two posts, the role of congestion is more interesting.

Example 4: $m = 2, n = 8$. The three types of agents and their 8-prefixes allow a single top-n-fair assignment

$$\text{prefixes} \quad \begin{array}{ccc} & a & b \\ \alpha\alpha\alpha\alpha & 8 & 0 \\ \beta\beta & 4 & 4 \\ \gamma\gamma & 0 & 8 \end{array} \quad \text{top-fair assignmt } P: \quad \begin{array}{cc} a & b \\ \alpha\alpha\alpha\alpha & \beta\beta\gamma\gamma \end{array}$$

Assignment P is competitive if the β-s prefer (weakly) $(b, 4)$ to $(a, 4)$; it is not if $(a, 4) \succ_\beta (b, 4)$ holds for at least one β.

Next we introduce a natural family of utility functions that will facilitate the randomisation exercise in section 3.2. Fixing $\lambda = (\lambda_a)_{a \in A}$ in $\Delta^{\mathbb{N}}(A; n)$ we define the

[4] For a simple example $A = \{a, b\}$, all but agent 1 refuse a and agent 1 is indifferent between $(a, 1)$ and (b, n); the two assignments $(S_a = \{1\}, S_b = N \diagdown \{1\})$ and $(S_a = \varnothing, S_b = N)$ are competitive but the former Pareto dominates the latter.

slack utility u^λ:

$$u^\lambda(a,s) = \lambda_a - s \text{ for all } (a,s) \in A \times [n] \tag{10}$$

It is linear in the congestion s and λ is the unique n-prefix of the corresponding preferences; all allocations (a, λ_a) are indifferent. In the next three examples everyone has slack utilities.

In Example 2 (section 2.2) it is easy to check that the single top-18-fair assignment is competitive as well.

In the next example $m = 3, n = 6$ and the agents are of three types. posts and six agents of three types α, β, γ. The λ-s and the three top-6-fair assignments (ignoring in the last one permutations of the α-s) are as follows

$$\text{prefixes} \quad \begin{array}{cccc} & a & b & c \\ \alpha\alpha & 3 & 2 & 1 \\ \beta\beta & 1 & 3 & 2 \\ \gamma\gamma & 2 & 1 & 3 \end{array} \quad \text{top-fair assignmts} \quad \begin{array}{ccc} a & b & c \\ \alpha\alpha & \beta\beta & \gamma\gamma \\ \gamma\gamma & \alpha\alpha & \beta\beta \\ \alpha\gamma & \alpha\beta & \beta\gamma \end{array} \tag{11}$$

We use property (3) to check $\mathcal{C}(\lambda) = \{(2,2,2)\}$, then we see that these are the only top-fair assignments. They all are competitive with identical utility profile $(0,0,0)$.

Finally we show that there is no competitive assignment in Example 3 (section 2.2) with slack utilities. By statement *iii*) above we only need to check this for the top-12-fair assignments: in the first one of table (4) all agents want to move to b; in the second one all agents at a or c envy those at b; and the reverse holds in the third one. It is easy to check similarly that none of the asymmetric congestion profiles in $\mathcal{C}(\lambda)$ is compatible with a competitive assignment.

3.1.1 free mobility equilibria and competitiveness

A competitive assignment is a FM equilibrium outcome: this is clear from Definition 2. The converse is not true because at least one FM equilibrium exists in every problem (section 2.2.1). Also, if a competitive assignment exists, it may be Pareto superior to other FM equilibrium outcomes.

Example 5: $m = 3, n = 6$. Agents of the same type have identical preferences described in the table, together with the two top-6-fair assignments

prefixes:

	a	b	c
$\alpha\alpha$	3	3	0
$\beta\beta$	3	0	3
$\gamma\gamma$	0	3	3

$(a,2) \succ_\alpha (b,2) \succ_\alpha (a,3)$
$(c,2) \succ_\beta (a,2) \succ_\beta (c,3)$
$(b,2) \succ_\gamma (c,2) \succ_\gamma (b,3)$

	a	b	c
P^1	$\alpha\alpha$	$\gamma\gamma$	$\beta\beta$
P^2	$\beta\beta$	$\alpha\alpha$	$\gamma\gamma$

Both P^1 and P^2 are FM equilibrium outcomes. Only P^1 is competitive, and it Pareto dominates P^2.

With more work we can also find problems where two FM equilibrium assignments can have quite different congestion and welfare profiles, in contrast to the consequences of competitiveness in Proposition 1 and Theorem 1 below. We give an example in section 6.1 with $m = 5$, $n = 26$.

3.2 the cardinal model: fractional competitiveness

Turning to a randomised assignment model will make the competitive approach operational by selecting in every problem a unique competitive congestion implemented by a lottery over approximately competitive deterministic assignments.

We start by endowing each agent i with a cardinal vNM utility $u_i(a, s)$ over $A \times [n]$. We extend each u_i to a function $u_i(a, \sigma)$ over $A \times [1, n]$ by linear interpolation between $\lfloor \sigma \rfloor$ and $\lceil \sigma \rceil$, the two rounded up and down values of σ:

$$u_i(a, \sigma) = \frac{\lceil \sigma \rceil - \sigma}{\lceil \sigma \rceil - \lfloor \sigma \rfloor} u_i(a, \lfloor \sigma \rfloor) + \frac{\sigma - \lfloor \sigma \rfloor}{\lceil \sigma \rceil - \lfloor \sigma \rfloor} u_i(a, \lceil \sigma \rceil) \text{ for all non-integer } \sigma \quad (12)$$

So $u_i(a, \sigma)$ is continuous and strictly decreasing in $\sigma \in [1, n]$.

For a finite set Z the notation $\Delta^{\mathbb{R}}(Z; y)$ is the simplex with non negative real coordinates in Z and sum the real number y. The key variable of a random assignment is the expected congestion σ_a at each post a: the congestion profile $\sigma = (\sigma_a)_{a \in A}$ varies in $\Delta^{\mathbb{R}}(A; n)$. Note that $\sigma_a < 1$ is possible: in this case the definitions below make sure that each agent perceives a congestion of 1 at a, so $u_i(a, \sigma)$ is *not* defined if $\sigma \in [0, 1[$. Moreover in Definition 4 describing the implementation of a random congestion σ by a lottery over deterministic assignments, the deterministic congestion at σ_a takes only the values $\lfloor \sigma_a \rfloor$ and $\lceil \sigma_a \rceil$, so that equation (12) is indeed the relevant vNM expected utility at (a, σ).

As in the proof of Proposition 2 we define agent i's *competitive demand* at the expected congestion profile $\sigma \in \Delta^{\mathbb{R}}(A;n)$:

$$D(u_i, \sigma) = \arg\max_{x \in A} u_i(x, \sigma_x \vee 1) \tag{13}$$

This demand ignores the effect of agent i's own presence at post x but correctly anticipates that the ex post congestion will round σ_x up or down. Our key definition comes next.

Definition 3: *In the problem (A, N, u) the fractional congestion profile $\sigma \in \Delta^{\mathbb{R}}(A;n)$ is competitive (Comp) iff*

$$\sigma \in \sum_{i \in N} \Delta^{\mathbb{R}}[D(u_i, \sigma); 1] \tag{14}$$

Property (14) means that we can achieve the congestion profile σ^* by assigning randomly each agent, with a well chosen probability distribution, over the posts in her competitive demand. Writing this probability distribution as (π_{ia}) we obtain a semi-stochastic matrix $\Pi = [\pi_{ia}] \in [0,1]^{N \times A}$ (each row sums to 1) realising σ and with competitive support:

$$\sigma_a = \sum_{i \in N} \pi_{ia} \text{ for all } a \in A \tag{15}$$

$$\pi_{ia} > 0 \Longrightarrow a \in D(u_i, \sigma) \text{ for all } i \in N, a \in A \tag{16}$$

There may be more than one matrix Π realising a given congestion profile σ.

Lemma 2: *In every problem (A, N, u) there is a unique competitive congestion profile σ^c, except possibly at some posts where $\sigma^c_a \leq 1$,*[5] *and a unique competitive demand.*

Proof *Existence.* It follows from the Kakutani fixed point theorem applied to the convex compact valued correspondence $\Gamma(\sigma) = \sum_{i \in N} \Delta^{\mathbb{R}}(D(u_i, \sigma); 1)$ mapping $\Delta^{\mathbb{R}}(A;n)$ into itself. To check that the graph of Γ is closed (implying that Γ is upper-semi-continuous) we take a sequence $(\sigma^t, \tau^t)_{t=1}^{\infty}$ in $[\Delta^{\mathbb{R}}(A;n)]^2$ converging to (σ, τ) and s. t. $\tau^t = \sum_{i \in N} \tau^t_i$ and $\tau^t_i \in \Delta^{\mathbb{R}}[D(u_i, \sigma^t); 1]$. We can find a subsequence

[5] If σ, σ^* are both competitive and $\sigma^*_a \neq \sigma_a$ at some post a, then $\sigma^*_a, \sigma_a \leq 1$.

such that all sequences $\{\tau_i^t\}$ converge, and all sets $D(u_i, \sigma^t)$ are constant in t, so that $\tau \in \Gamma(\sigma)$ as desired.

Uniqueness. We adapt the proof of statement i) in Proposition 1. Assume σ, σ^* are two different competitive profiles with corresponding matrices Π and Π^* in (15), (16). We set $A^* = \{a \in A | \sigma_a^* \vee 1 > \sigma_a \vee 1\}$ and check, exactly like in the deterministic proof, that if an agent i is s. t. $D(u_i; \sigma^*)$ intersects A^* at some a, then $D(u_i; \sigma)$ must be a subset of A^*.

Therefore for any i s. t. $\sum_{a \in A^*} \pi_{ia}^* > 0$ we have $\sum_{a \in A^*} \pi_{ia}^* \leq 1 = \sum_{a \in A^*} \pi_{ia}$; summing over all agents this gives $\sum_{a \in A^*} \sigma_a^* = \sum_{i \in N, a \in A^*} \pi_{ia}^* \leq \sum_{i \in N, a \in A^*} \pi_{ia} = \sum_{a \in A^*} \sigma_a$, a contradiction of the assumption $\sigma_a^* > \sigma_a$ in A^*. We conclude that A^* is empty: $\sigma_a^* \vee 1 \leq \sigma_a \vee 1$ for all a. The opposite inequality holds by exchanging the roles of σ and σ^*, implying $\sigma_a^* \vee 1 = \sigma_a \vee 1$ and the desired conclusion that $\sigma_a^* \neq \sigma_a$ can only happen when both are at most 1.

Because an agent evaluates being assigned to a post with expected congestion at most 1 as being alone there (definition (13)) this implies $D(u_i, \sigma) = D(u_i, \sigma^*)$, therefore σ and σ^* generate the same competitive demand. ■

Definition 4: *In the problem (A, N, u) the list $(\{P^k\}_{k=1}^K, \mathcal{L})$ of K deterministic assignments P^k together with a lottery $\mathcal{L}$ over $[K]$, implements the competitive congestion profile σ^c if, first, the expected congestion over these K assignments is σ^c: $\mathbb{E}_{\mathcal{L}}(s_a^k) = \sigma_a^c$ for all a; and second for all $k \in [K]$, all i, all a and all k we have:*

$$i \in S_a^k \Longrightarrow a \in D(u_i, \sigma^c) \text{ and } s_a^k = \lfloor \sigma_a^c \rfloor \text{ or } \lceil \sigma_a^c \rceil \tag{17}$$

Property (17) says that anyone of the assignments P^k is ex post fair in the sense that each agent is at a post in his competitive demand (based on the expected congestion σ^c) and the actual congestion at P^k is an integer rounding of the latter.

Lemma 3: *In every problem (A, N, u) we can implement the competitive congestion profile σ^c by (typically several) lists $(\{P^k\}_{k=1}^K, \mathcal{L})$ as in Definition 4.*

Corollary: *If σ^c is integer-valued, each assignment P^k is competitive and implements σ^c.*

Proof Define the set $\mathcal{S}$ of semi-stochastic matrices Π s.t. for all a and i: $\{\pi_{ia} > 0 \Longrightarrow a \in D(u_i, \sigma^c)\}$ and $\lfloor \sigma_a^c \rfloor \leq \sum_{i \in N} \pi_{ia} \leq \lceil \sigma_a^c \rceil$.This set is a convex compact polytope, non empty as it contains any matrix Π^c realising σ^c. We claim that each extreme point Π^k of $\mathcal{S}$ is deterministic, i. e., its entries are all integers or zero; so Π^k is a deterministic assignment P^k meeting (17) and Π^c is a convex combination of such extreme points: this gives us the desired collection P^k and lottery $\mathcal{L}$.

We prove the claim by contradiction. Pick an extreme point Π of $\mathcal{S}$ and associate to Π the bipartite graph G on $N \times A$ containing the edge ia iff $\pi_{ia} > 0$. Extract from G the subgraph G_0 of its fractional entries ia, i. e., $0 < \pi_{ia} < 1$, and let F be the set of posts a s.t. $\sum_{i \in N} \pi_{ia}$ is fractional (not an integer or zero). If F is non empty it contains some post a: at least one edge ia is in G_0; then at least one other edge ib is in G_0 (Π is semi-stochastic at i); if $b \in F$ then we can add a small ε to π_{ia} and take away ε from π_{ib} without leaving $\mathcal{S}$, or vice versa: this contradicts the extremality of Π, so b is not in F after all. But then there is another edge jb in G_0 because $\sum_{i \in N} \pi_{ib}$ is an integer, and again there is some new edge jc in G_0; if $c = a$ or $c \in F$ we can as above perturb a little the entries of Π in two opposite ways and reach a contradiction. This construction must stop at a new post in F or cycle back to a: in both cases we can perturb the path or cycle in opposite directions. The claim is proved and the proof is complete. ■

Our main result refines the properties of the deterministic assignments P^k selected ex post to implement the competitive congestion: any two such assignments yield approximately identical utilities; and each P^k shares (approximately) the properties of deterministic competitive assignments in Proposition 1. We use for agent i's approximation parameter her worst utility loss from one additional unit of congestion $\delta_i = \max_{(a,s) \in A \times [n]} \{u_i(a, s) - u_i(a, s + 1)\}$. For instance $\delta_i = 1$ for a slack utility (10).

Theorem 1 *Fix a problem* (A, N, u)*, a list* $(\{P^k\}_{k=1}^K, \mathcal{L})$ *implementing the competitive congestion* σ^c (Definition 4) *and write* $U^k = (U_i^k)_{i \in N}$ *the utility profile at assignment* P^k*.*
i) Each P^k *is top-n-fair up to at most one unit of congestion: for all* $i \in N$ *there is a n-prefix* λ_i *of* u_i *such that* $i \in S_a^k \Longrightarrow s_a^k \leq \lambda_{ia} + 1$*.*
ii) Utilities at two assignments P^k, P^ℓ *differ by at most* $2\delta_i$*:* $|U_i^k - U_i^\ell| \leq 2\delta_i$ *for all* i *and all* $k, \ell \in [K]$*.*
iii) Each P^k *is* $2\delta_i$*-competitive:* $U_i^k \geq u_i(x, s_x^k \vee 1) - 2\delta_i$ *for all* i *and all* a*.*

iv) Each assignment P^k is $(2\delta_i + \varepsilon)$-efficient for any $\varepsilon > 0$: no assignment $Q \in \mathcal{P}(A, N)$ Pareto dominates the profile $(U_i^k + 2\delta_i + \varepsilon)_{i \in N}$.

If in addition $\sigma_a^c \geq 1$ for all a, we can say more. Let U_i^c be the value $u_i(a, \sigma_a^c)$ common to all $a \in D(u_i, \sigma^c)$, then the profile $(U_i^c)_{i \in N}$ is efficient[6]*; finally we have $U_i^k > U_i^c - \delta_i$ for all i and all $k \in [K]$.*

Proof *Statement i)* Fix P^k and write its congestion profile simply as $(s_x)_{x \in A}$, then fix an agent i and her allocation (a, s_a) at P^k; finally $[\sigma^c]$ is the support of the competitive congestion, containing x iff $\sigma_x^c > 0$.

Suppose first $s_a \geq 2$. By (17) we have $s_a - 1 \leq \lfloor \sigma_a^c \rfloor$ therefore $\sigma_a^c \geq 1$. We apply repeatedly the monotonicity of $u_i(x, s)$ in s and property (17).
For all $x \in [\sigma^c]$: $u_i(a, s_a - 1) \geq u_i(a, \lfloor \sigma_a^c \rfloor) \geq u_i(a, \sigma_a^c) \geq u_i(x, \sigma_x^c \vee 1) \geq u_i(x, \lceil \sigma_x^c \rceil)$ .where the last inequality follows from $\sigma_x^c > 0$. These inequalities imply that agent i prefers to $(a, s_a - 1)$: at most $\lfloor \sigma_a^c \rfloor - 1$ less congested allocations at post a; at most $\lceil \sigma_x^c \rceil - 1 \leq \lfloor \sigma_x^c \rfloor$ allocations at post x if $\sigma_x^c > 0$; and if $\sigma_y^c = 0$ (17) gives $u_i(a, s_a) \geq u_i(y, 1)$ so no allocation at y improves (a, s_a) or $(a, s_a - 1)$ for our agent i. We see that the number of allocations improving $(a, s_a - 1)$ is at most $(\sum_{x \in A} \lfloor \sigma_x^c \rfloor) - 1 \leq \sum_{x \in A} \sigma_x^c - 1 = n - 1$. As in section 2.2.1 we conclude that $(a, s_a - 1)$ is top-n-fair, and we are done.

In the case $s_a = 1$ the argument is simpler. The allocation $(a, 1)$ is the best at post a; at any other post x the inequality $u_i(a, 1) \geq u_i(x, \sigma_x^c \vee 1)$ implies it can be improved by at most $\lfloor \sigma_x^c \rfloor - 1$ less congested allocations (or zero if $\sigma_x^c < 1$). Then the allocation (a, s_a) itself is top-n-fair.
Statement ii) Fix P^k, P^ℓ, a, b and $i \in S_a^k \cap S_b^\ell$. By (17) a, b are both in $D(u_i, \sigma^c)$ hence $u_i(a, \sigma_a^c \vee 1) = u_i(b, \sigma_b^c \vee 1)$. By (17) again and the definition of δ_i we have $|u_i(a, s_a^k) - u_i(a, \sigma_a^c \vee 1)| \leq \delta_i$ and $|u_i(b, s_b^\ell) - u_i(b, \sigma_b^c \vee 1)| \leq \delta_i$, so the desired inequality follows.
Statement iii) We fix $i \in S_a^k$ and x then combine three inequalities: $a \in D(u_i, \sigma^c)$ gives $u_i(a, \sigma_a^c \vee 1) \geq u_i(x, \sigma_x^c \vee 1)$; next $|u_i(a, s_a^k) - u_i(a, \sigma_a^c \vee 1)| \leq \delta_i$ follows from $|s_a^k - \sigma_a^c| < 1$ ((17)) and the definition of δ_i; similarly $|u_i(x, s_x^k) - u_i(a, \sigma_x^c \vee 1)| \leq \delta_i$.
Statement iv) For the first part of the statement we fix P^k and $Q = (T_x)_{x \in A}$ in

[6] Not dominated by the utility profile of any *assignment* $Q \in \mathcal{P}(A, N)$.

$\mathcal{P}(A, N)$. There is at least one post b in the support of Q ($T_b \neq \varnothing$) such that $s_b^k \leq t_b$: otherwise $t_b < s_b^k$ for each b s. t. $t_b \geq 1$, which contradicts $\sum_A t_b = \sum_A s_b^k = n$. Pick such a post b, some $i \in T_b$, and let a be the post assigned to i by P^k ($a = b$ is possible). By statement $iii)$ we have $u_i(a, s_a^k) \geq u_i(b, s_b^k \vee 1) - 2\delta_i$ and $u_i(b, s_b^k \vee 1) \geq u_i(b, t_b)$ by our choice of b. So Q does not improve P^k by more than $2\delta_i$ for all agents in T_b.

For the second part we have $\sigma_a^c \geq 1$ for all a. Assume that the utility profile of a deterministic assignment $Q = (T_x)_{x \in A}$ Pareto dominates $(U_i^c)_{i \in N}$. For any post a in the support of Q and agent $i \in T_a$, the definition of the competitive demand ((13)) and the choice of Q imply $u_i(a, \sigma_a^c) \leq U_i^c \leq u_i(a, t_a)$ and in turn $t_a \leq \sigma_a^c$ for all a. As $\sum_a t_a = \sum_a \sigma_a^c$ we conclude $t = \sigma^c$ and the inequalities above are equalities, so Q's utility profile is exactly $(U_i^c)_{i \in N}$, and Q does not dominate $(U_i^c)_{i \in N}$ after all.

Next for any P^k, any i and a such that $i \in S_a^k$, Definition 4 implies $s_a^k \leq \lceil \sigma_a^c \rceil < \sigma_a^c + 1$ hence $u_i(a, s_a^k) > u_i(a, \sigma_a^c + 1) \geq u_i(a, \sigma_a^c) - \delta_i = U_i^c - \delta_i$ as desired. ■

3.2.1 computing competitive assignments: examples

Example 4 $m = 2, n = 8$ (section 3.1) To the 8-prefixes of the agents we add the critical part in the cardinal utilities of the two β agents:

	a	b
$\alpha\alpha\alpha\alpha$	8	0
$\beta\beta$	4	4
$\gamma\gamma$	0	8

	$(a,5)$	$(b,4)$	$(a,4)$	$(b,3)$
u_{β_1}	0	0	1	2
u_{β_2}	-1	0	1	2

The β-s are envious at the single top-n-fair assignment $P^1 = \begin{array}{cc} a & b \\ \alpha\alpha\alpha\alpha & \beta\beta\gamma\gamma \end{array}$ because $(a, 4) \succ_\beta (b, 4)$ where β_1 suffers comparatively less than β_2 at $(a, 5)$.

The competitive fractional congestion must load post a more than post b: $\sigma^c = (4 + x, 4 - x)$ where $x \geq 0$; then some deterministic assignment implementing σ^c will violate top-8-fairness. By statement $i)$ in the Theorem we have $x \leq 1$, so the demand of the α-s and γ-s will not change; to get $D(u_{\beta_1}, \sigma) = \{a\}$, $D(u_{\beta_2}, \sigma) = \{a, b\}$ the only choice is $x = \frac{1}{4}$.

So $\sigma^c = (4\frac{1}{4}, 3\frac{3}{4})$ meets (14).[7] The lottery $\frac{3}{4}P^1 + \frac{1}{4}P^2$ with $P^2 = \begin{array}{cc} a & b \\ \alpha\alpha\alpha\alpha\beta_2 & \beta_1\gamma\gamma \end{array}$ implements it.

Note that for agent β_2, conditional on being at post a the expected congestion there is 5, not $4\frac{1}{4}$: reasoning competitively β_2 accepts both posts at σ^c without taking into account the impact of her own assignment to a on the price/congestion at a; by property (17) this discrepancy is at most one unit of congestion, irrespective of the size of n and m.

Comparing the competitive lottery $\frac{3}{4}P^1 + \frac{1}{4}P^2$ and the top-n-fair assignment P^1 agent β_1 and the γ-s prefer the lottery while β_2 and the α-s have the opposite preference.

Example 3 $m = 3, n = 12$ (section 2.2) We have two types of agents with the 12-prefixes $\begin{array}{cccc} & a & b & c \\ \alpha\alpha\alpha\alpha\alpha\alpha & 6 & 4 & 2 \\ \beta\beta\beta\beta\beta\beta & 2 & 4 & 6 \end{array}$. The slack utility (10) applies to each agent: their vNM extension to real valued congestion (12) still is $u^{\lambda_i}(a, \sigma_a) = \lambda_{ia} - \sigma_a$.

We checked in section 3.1 that no deterministic assignment is competitive. The unique competitive fractional congestion respects the symmetries of the problem (this is a general property following its definition by the fixed point property (14)), so we look for $\sigma^c = (x, y, x)$ with $2x + y = 12$. The only choice generating the demands $D(u^{\lambda_\alpha}, \sigma) = \{a, b\}$, $D(u^{\lambda_\beta}, \sigma) = \{b, c\}$ is $\sigma^c = (4\frac{2}{3}, 2\frac{2}{3}, 4\frac{2}{3})$. To implement σ^c we must combine top-12-fair deterministic assignments where $s_a, s_c \in \{4, 5\}$ and $s_b \in \{2, 3\}$, which leaves exactly three choices

	a	b	c
Q^1	$\alpha\alpha\alpha\alpha\alpha$	$\alpha\beta$	$\beta\beta\beta\beta\beta$
Q^2	$\alpha\alpha\alpha\alpha\alpha$	$\alpha\beta\beta$	$\beta\beta\beta\beta$
Q^3	$\alpha\alpha\alpha\alpha$	$\alpha\alpha\beta$	$\beta\beta\beta\beta\beta$

The lottery $L^c = \frac{1}{3}Q^1 + \frac{1}{3}Q^2 + \frac{1}{3}Q^3$ (plus a uniform mixing of the α-s and of the β-s in their respective roles) is our, in this case unique, implementation of σ^c.

[7] Whereas at $x' = \frac{1}{3}$ we have $D(u_{\beta_1}, \sigma) = \{a, b\}$, $D(u_{\beta_2}, \sigma) = \{b\}$ and (14) fails.

The expected total utility of the α-s is $\frac{1}{3}(7+6+10)$ so each agent's expected utility is 1.28. But the symmetric and top-12-fair assignment $Q^4 = \begin{matrix} a & b & c \\ \alpha\alpha\alpha\alpha & \alpha\alpha\beta\beta & \beta\beta\beta\beta \end{matrix}$ collects more expected utility: after the uniform mixing inside each type, everyone gets 1.33. The trade-off is between gathering more surplus at Q^4 but allowing $s_b = 4$ so that the agents assigned to b have an ex post envy at the level of 2 units of congestion, versus losing a little surplus at the competitive lottery L^c while generating ex post envy only at the level of 1 unit of congestion.

3.2.2 failure of the ordinal randomisation

Here we assume that the profile of ordinal preferences is the only welfare information available. As in most of the literature on random assignment (see section 1.1) we use stochastic dominance between random congested assignments to define both efficiency and envy-freeness, and look for fractional assignments combining these two properties and (approximate) top-n-fairness. This turns out to be easy but much too permissive.

Fix a problem $(A, N, u_i, i \in N)$ and an *arbitrary* deterministic congestion profile $s \in \Delta^{\mathbb{N}}(A; n)$. For each allocation (a, s_a) create s_a copies of an item labeled $\widetilde{a}$, for a total of n new items, then define a *standard* assignment problem with n agents labeled $\widetilde{i}$ who compare each copy of $\widetilde{a}$ and of $\widetilde{b}$ as i compares (a, s_a) and (b, s_b). In this "decongested" problem we run the version of the PS algorithm allowing for indifferences (see **[21], [4]**), and obtain a random assignment $\widetilde{\Pi} \in [0,1]^{N \times \widetilde{A}}$ (a fully stochastic matrix) SD-efficient and SD-envy-free. Back in the original congested problem the matrix $\widetilde{\Pi}$ defines a random assignment implementing congestion s where for all agents only the allocations $(a, s_a), a \in A$, can have positive probability; this assignment is clearly SD-efficient and SD-envy-free. Note that the decongestion method works just as well if we start from a fractional congestion profile σ instead of s, provided the corresponding expected allocations (x, σ_x) are SD-comparable with ordinal preferences (see an example below).

In Example 3 we illustrate the purely ordinal approach with the slack preferences between deterministic allocations, instead of cardinal slack utilities in section 3.2.1. Two natural starting congestion profiles are the two symmetric top-8-fair congestion profiles: $s^1 = (5,2,5)$ at assignment Q^1 and $s^4 = (4,4,4)$ at Q^4.

Choose s^1: all agents eat first the two units of $\widetilde{b}$ before the α-s turn to $\widetilde{a}$ and the β-s to $\widetilde{c}$; the random assignment $\frac{5}{6}(a,5) + \frac{1}{6}(b,2)$ to the α-s, $\frac{5}{6}(c,5) + \frac{1}{6}(b,2)$ to the β-s, is implemented by mixing uniformly agents of the same type in Q^1.

Choose s^4: the α-s eat first the four $\widetilde{a}$ before sharing the $\widetilde{b}$-s with the β-s; uniform mixing by types produces the random assignment $\frac{2}{3}(a,4) + \frac{1}{3}(b,4)$ for the α-s, $\frac{2}{3}(c,4) + \frac{1}{3}(b,4)$ for the β-s.

Observing that the number of envious agents is minimal at assignments Q^2, Q^3 (only two β-s in Q^2; but ten in Q^1 and four in Q^4) it also makes sense to start from the expected congestion $(4\frac{1}{2}, 3, 4\frac{1}{2})$ at $\frac{1}{2}Q^2 + \frac{1}{2}Q^3$. Indeed if $(a, 4\frac{1}{2})$ is implemented as $\frac{1}{2}(a,5) + \frac{1}{2}(a,4)$, for an α agent it is stochastically dominated by $(b,3)$ so in the decongested problem with $4\frac{1}{2}$ units of $\widetilde{a}$ and of $\widetilde{c}$ and 3 units of $\widetilde{b}$ the PS algorithm is unambiguous.[8] Then all agents eat first three units of $\widetilde{b}$ and the random assignment $\frac{3}{8}(a,5) + \frac{3}{8}(a,4) + \frac{1}{4}(b,3)$ for the α-s, and similarly for the β-s, is implemented by fairly mixing first Q^2 and Q^3, then the agents of the same type.

We end up with an embarrassment of riches, and no systematic way to select between these three plausible randomised solutions.

In appendix 6.2 we illustrate another problem of the ordinal approach: it may generate unbounded ex post envy, in the sense that its implementation may force a positive probability on deterministic assignments where an agent's allocation is, in her view, many ranks below that of another agent.

4 weighted congestion

Now each agent i brings the amount w_i of congestion to her assigned post: w_i is a *strictly positive* real number. The total congestion is $W = \sum_{i \in N} w_i$. For any subset S of N we use the notation $w_S = \sum_{j \in S} w_j$.

We discuss first the deterministic model with ordinal preferences (sections 4.1, 4.2) and turn to randomisation and cardinal utilities in section 4.3.

[8]But ordinal preferences are not enough to SD-compare $\frac{1}{2}(a,6) + \frac{1}{2}(a,3)$ and $(b,3)$.

4.1 ex ante fairness: the top-$\frac{1}{m}$-guarantee

Ex ante fairness for agent i only depends upon her own weight w_i, the total congestion W and the set of posts A. It is independent of the number of other agents filling the congestion $W - w_i$. Thus the weighted model is not a direct generalisation of the anonymous congestion one: an agent must allow for any finite split of the total weight $W - w_i$ of other participants, which seriously enlarges the set of adversarial situations to take into account.

We write $\widetilde{\mathcal{F}}_i = A \times [w_i, W]$ the set of agent i's *potential* allocations at the ex ante stage and endow each agent i with continuous preferences $\widetilde{\succsim}_i$ over $\widetilde{\mathcal{F}}_i$. The size of $\widetilde{\mathcal{F}}_i$ is $m(W - w_i)$: we show below that agent i can guarantee one of her top $\frac{1}{m}$ allocations in $\widetilde{\mathcal{F}}_i$, i. e., in a subset of size $(W - w_i)$. This is exactly like in the anonymous model where out of $m \times n$ allocations in total, agent i's n-prefixes contain the top $\frac{1}{m}$ quantile of her welfare levels.

The description of these $(W-w_i)$-prefixes is more subtle than under anonymous congestion, but on the bright side this prefix is unique for each preference $\widetilde{\succsim}_i$. With the notation $[[z]]$ for the number of strictly positive coordinates of $z \in \mathbb{R}^A_+$ a $(W-w_i)$-prefix is described by a vector $\lambda_i = (\lambda_{ia})_{a \in A}$ such that

$$\text{for each } a : \lambda_{ia} = 0 \text{ or } w_i \leq \lambda_{ia} \leq W, \text{ and } \sum_{a \in A} \lambda_{ia} = W + ([[\lambda_i]] - 1)w_i \qquad (18)$$

This defines the subset $\mathcal{G}(A; w_i; W)$ of $(\{0\} \cup [w_i, W])^A$.

We interpret equation (18) as follows. If $\lambda_{ia} \in [w_i, W]$ all allocations in $\{a\} \times [\lambda_{ia} - w_i]$ are allowed; in particular $\lambda_{ia} = w_i$ means that agent i can be at post a only if she is alone there. If $\lambda_{ia} = 0$ agent i will not be assigned to post a: notice that changing λ_{ia} from 0 to w_i adds w_i to both sides of (18), therefore ruling post a out entirely or accepting it only if she is alone there has no impact on the feasibility constraint $\mathcal{G}(A; w_i; W)$. For instance if $\lambda_{ia} = W$ our agent can still choose to accept any subset of the other posts provided she is alone there.

Let B be the (strict and possibly empty) subset B of the posts such that $\lambda_i = 0$; the total size of the allocations in $(A \diagdown B) \times [w_i, W]$ that λ_i allows is

$$\sum_{a \in A \diagdown B} (\lambda_{ia} - w_i) = \sum_{a \in A} \lambda_{ia} - (m - |B|)w_i = W - w_i$$

so λ_i can describe precisely the $(W - w_i)$-prefix of $\succeq_i$,[9] as announced earlier.

Definition 5 *In the problem* $(A, N, \widetilde{\succeq}, w)$ *where* $\lambda_i \in \mathcal{G}(A; w_i; W)$ *is the* $(W - w_i)$*-prefix of* $\widetilde{\succeq}_i$*, an assignment* P *is top-*$\frac{1}{m}$*-fair iff*

$$\sum_{j \in S_a} w_j \leq \lambda_{ia} \text{ for all } a \in A \text{ and all } i \in S_a$$

Proposition 2 *In any problem* $(A, N, \widetilde{\succeq}, w)$ *there exists at least one top-*$\frac{1}{m}$*-fair assignment* P*.*

Maximality property: Fix A, W*, an agent* i^* *with weight* w_{i^*} *and an arbitrary* $(W - w_{i^*})$*-prefix* $\lambda_{i^*} \in \mathcal{G}(A; w_{i^*}; W)$*. Then for any* a s. t. $\lambda_{i^*a} > w_{i^*}$ *there is a set of agents* M *not containing* i^* *with weights* w_i *s. t.* $\sum_{i \in M} w_i = W - w_{i^*}$ *and a prefix* $\lambda_i \in \mathcal{G}(A; w_i; W)$ *for each* $i \in M$*, s. t. in any top-n-fair assignment of the* $M \cup \{i^*\}$ *problem agent* i^* *is at a post* a *where the congestion is arbitrarily close to* λ_{i^*a}*.*

Proof *Existence.* As in Lemma 1 we combine a greedy algorithm and an induction on the number of agents. We can clearly get rid of the "inactive" posts s. t. $\lambda_{ia} = 0$ for all i, so we still write A for the set of active posts. We pick any post a and construct a set $S \subseteq A$ s. t.

$$\forall i \in S : \lambda_{ia} \geq w_S \text{ and } \forall j \in N \diagdown S : \lambda_{ja} < w_S + w_j \tag{19}$$

Label the agents from 1 to n so that $\lambda_{ia} \geq \lambda_{(i+1)a}$ for all $i \in [n-1]$ where $\lambda_{1a} > 0$ because a is active. For any two disjoints subsets S, T in A we say that S rejects T if $\lambda_{ia} < w_S + w_T$ for some $i \in S$; otherwise we say that S accepts T. Note that if all labels in S are (weakly) smaller than all in T and S rejects T, then T rejects S as well; and S accepts T if T accepts S. We construct S recursively: in each step we either find S or add one agent to the provisional set.

step 1. If for all $j \geq 2$ the set $\{j\}$ rejects $\{1\}$ then $S = \{1\}$ meets (19) and we are done. Otherwise we pick the smallest label $\ell^1 \geq 2$ accepting $\{1\}$, which implies that $\{1\}$ accepts $\{\ell^1\}$ as well, and we form the provisional set $S^1 = \{1, \ell^1\}$. If $\ell^1 = n$ we are done by choosing S^1 so going into step 2 we have $\ell^1 < n$.

[9] We let the reader check that this set is uniquely defined.

step $k+1$. The subset S has not yet been found therefore ℓ^k, the latest pick in S^k, is smaller than n. By construction $\lambda_{\ell^k a} \geq w_{S^{k-1}} + w_{\ell^k} = w_{S^k}$ so $\lambda_{ia} \geq w_{S^k}$ for all $i \in S^k$. Moreover all agents $j < \ell^k$ outside S^k have rejected some earlier $S^{k'}$, so they also reject the larger set S^k.

If all agents $j > \ell^k$ reject S^k as well we are done by choosing S^k. Otherwise we pick the smallest label ℓ^{k+1} after ℓ^k s.t. $\{\ell^{k+1}\}$ accepts S^k: this implies that S^k accepts $\{\ell^{k+1}\}$ as well (recall $\lambda_{\ell^k a} \geq \lambda_{\ell^{k+1} a}$) so we set $S^{k+1} = S^k \cup \{\ell^{k+1}\}$ and we have $\lambda_{ia} \geq w_{S^{k+1}}$ for all $i \in S^{k+1}$. We are done if $\ell^{k+1} = n$ otherwise we go to the next step. When this process stops we have found $\widetilde{S}$ meeting (19).

We assign $\widetilde{S}$ to post a, and consider the residual problem in $\widetilde{A} = A \diagdown \{a\}$, $\widetilde{N} = N \diagdown \widetilde{S}$ with total congestion $\widetilde{W} = W - w_S$. For each $j \in \widetilde{A}$ such that $\lambda_{ja} > 0$ inequality (19) and equation (18) imply

$$\lambda_{j\widetilde{A}} > \lambda_{jA} - (w_S + w_j) = W - w_S + ([[\lambda_j]] - 2)w_j = \widetilde{W} + ([[\widetilde{\lambda}_j]] - 1)w_j$$

where in $\widetilde{\lambda}_j$ we drop λ_{ja}. For each $j \in \widetilde{A}$ such that $\lambda_{ja} = 0$ inequality (19) has no bite and equation (18) gives

$$\lambda_{j\widetilde{A}} = W + ([[\lambda_j]] - 1)w_j > \widetilde{W} + ([[\lambda_j]] - 1)w_j = \widetilde{W} + ([[\widetilde{\lambda}_j]] - 1)w_j$$

The induction argument shows that in the reduced problem there is a top-$\frac{1}{m-1}$-fair assignment of $N \diagdown S$ to $\widetilde{A}$.
Maximality We fix W, i^*, the $(W - w_{i^*})$-prefix λ_{i^*} and a as in the statement. The set M contains one agent i_b for each post b s. t. $\lambda_{i^*b} \geq w_{i^*}$, in particular one i_a. The size of M is $[[\lambda_{i^*}]]$. The weights are $w_{i_b} = \lambda_{i^*b} - w_{i^*} + \varepsilon$ if $b \in M \diagdown a$ and $w_{i_a} = \lambda_{i^*a} - w_{i^*} - ([[\lambda_{i^*}]] - 1)\varepsilon$, where $\varepsilon > 0$ is small enough that $w_{i_a} > 0$. Clearly $w_{i^*} + \sum_{b \in M} w_{i_b} = W$. Suppose now that each i_b (including i_a) is single-minded on post b: $\lambda_{i_b b} = W$, $\lambda_{i_b x} = 0$ for $x \neq b$. Consider a top-$\frac{1}{m}$-fair assignment of $M \cup \{i^*\}$ to A: for each i_b other than i_a the inequality $w_{i_b} + w_{i^*} > \lambda_{i^*b}$ implies that i_b is alone at b, therefore i^* share a with i_a where the congestion $\lambda_{i^*a} - ([[\lambda_{i^*}]] - 1)\varepsilon$ is arbitrarily close to λ_{i^*a}. ■

Finding the top-$\frac{1}{m}$-fair assignments is more difficult with weighted rather than unweighted congestion. Already with two posts, multiple top-n-fair congestion profiles is quite normal, contrasting with the anonymous congestion case where there is only one (Example 1 section 2).

Example 6: *$m = 2, n = 3$, and $W = 10$.* The weights and $(W - w_i)$-prefixes are

	a	b	w
α	10	6	6
β	6	6	2
γ	8	4	2

For instance α's feasible set $\widetilde{\mathcal{F}}_\alpha$ is $\{a\} \times [6, 10] \cup \{b\} \times [6, 10]$ and his 4-prefix contains the allocations $\{a\} \times [6, 10] \cup \{(b, 6)\}$.

We show the three top-$\frac{1}{2}$-fair assignments and the corresponding utility profiles when the preferences are represented by the slack utilities $u_i(a, w_{s_a}) = \lambda_{ia} - w_{s_a}$

top-$\frac{1}{2}$-fair:

	a	b
P^1	α	$\beta\gamma$
P^2	$\alpha\gamma$	β
P^3	$\beta\gamma$	α

utilities:

	α	β	γ
P^1	4	2	0
P^2	2	4	0
P^3	0	2	4

The point is that the congestion and utility profiles are very different, in contrast with two post problems under anonymous congestion (Example 1 section 2.2).

Remark 3: Free mobility equilibrium and top-$\frac{1}{m}$-fairness. An equilibrium assignment of the FM game is top-$\frac{1}{m}$-fair, just like in the the anonymous congestion model (section 2.2.1). To check this fix such an assignment P and let (a, w_{S_a}) be agent i's equilibrium allocation. To (a, w_{S_a}) our agent prefers the allocations $\{a\} \times [w_i, w_{S_a}]$ and possibly those in $\{x\} \times [w_i, w_{S_x} + w_i]$: the total length of those sets is exactly $W - w_i$, which proves the claim.

But we know that the free mobility game may have no Nash equilibrium when congestion is weighted,[10] which makes the FM game much less relevant in this case.

[10] See an example in section 8 of [24].

4.2 deterministic competitiveness

At the ex post stage the number n of agents and their weights are known, so the range of agent i's feasible allocations is the finite set

$$\mathcal{F}_i = A \times \{w_S : \text{for some } S \text{ s.t. } i \in S \subseteq N\}$$

We write i's preferences over $\mathcal{F}_i$ as $\succeq_i$ (of course consistent with $\widetilde{\succeq}_i$ over $\widetilde{\mathcal{F}}_i$); as before the only restriction on $\succeq_i$ is to be strictly decreasing in congestion.

Definition 6 *Fix a problem* $(A, N, \succeq, w)$. *The assignment* P *is competitive (Comp) iff for all* $a, x \in A$ *and all* $i \in S_a$: $(a, w_{S_a}) \succeq_i (x, w_{S_x} \vee w_i)$.

The interpretation is the same as when congestion is anonymous, except for the fact that the "congestion price" of a post is now agent-specific. As in Proposition 1 a competitive assignment is efficient and top-fair, but the critical uniqueness property requires at least one of two qualifying properties.

Definition 7 *The assignment* P *crowded iff* $|S_a| \geq 2$ *or* $S_a = \varnothing$ *for all* a.
The ordinal preferences $\succeq_i$ *are semi-strict iff for any two posts* a, b *and any agent* i *and* S *s. t.* $i \in S \subseteq N$, *agent* i *is not indifferent between* (a, w_i) *and* (b, w_S).

Proposition 3: *Fix a problem* $(A, N, \succeq, w)$.
i) If all preferences $\succeq_i$ *are semi-strict then all competitive assignments have the same congestion and welfare profiles.*
The same is true if at least one competitive assignment is crowded.
ii) If all preferences are semi-strict a competitive assignment is efficient.
iii) A competitive assignment is top-$\frac{1}{m}$*-fair*

Proof It mimicks that of Proposition 1. Given the assignment $P = (S_x)_{x \in A}$ for simplicity we write the congestion w_{S_x} simplys as s_x. Now agent i's competitive demand at the congestion profile s is
$D(i; s) = \{a | (a, s_a \vee w_i) \succeq_i (x, s_x \vee w_i) \text{ for all } x \in A\}$, and P is competitive iff $a \in D(i; s)$ whenever $i \in S_a$.

Statement i) Unique competitive congestion. At first we do not assume in $(A, N, \succeq, w)$ either semi-strict preferences or a crowded competitive assignment.

Fix $P = (S_x)_{x \in A}$ and $P^* = (S^*_x)_{x \in A}$ both competitive and such that $s \neq s^*$. Partition A as $B^* \cup C \cup B$ where $B^* = \{a : s_a < s^*_a\}$, $B = \{a : s^*_a < s_a\}$,

$C = \{a : s^*_a = s_a\}$, and B, B^* are both non empty. We define in two equivalent ways the set $A^* = \{a \in B^* : s^*_a > w_i \text{ for all } i \in S^*_a\} = \{a \in B^* : |S^*_a| \geq 2\}$.

We claim first (as in the proof of Proposition 1) that if the demand $D(i; s^*)$ of some agent i intersects A^*, then $D(i; s) \subseteq A^*$. Suppose, to the contrary, $a \in D(i; s^*) \cap A^*$ while $b \in D(i; s) \cap A \diagdown A^*$. By definition of the demands and of B^* we have

$$(a, s^*_a) \succeq_i (b, s^*_b \vee w_i) \text{ and } (b, s_b \vee w_i) \succeq_i (a, s_a \vee w_i) \succ_i (a, s^*_a)$$

where the last relation is strict because $s_a \vee w_i < s^*_a$. Note that we cannot replace $s_b \vee w_i$ by s_b because i may not be in S_b.

If $b \notin B^*$ we have $s_b \vee w_i \geq s^*_b$ so we add $(b, s^*_b \vee w_i) \succeq_i (b, s_b \vee w_i)$ to the above preferences to get a contradiction. If $b \in B^*$ then $|S^*_b| = 1$ because $b \notin A^*$, therefore $s_b < s^*_b = w_i$ so that $(b, s_b \vee w_i) = (b, s^*_b \vee w_i)$ and we reach again a contradiction. The claim is proved. As in the proof of Proposition 1 we check next that A^* must be empty: every $i \in \cup_{a \in A^*} S^*_a$ is s. t. $D(i; s) \subseteq A^*$ therefore $\sum_{a \in A^*} s_a \geq \sum_{a \in A^*} s^*_a$, which contradicts $s^*_a > s_a$ in B^*. So for every $a \in B^*$ we have $S^*_a = \{i\}$. Exchanging the roles of P and P^* we see that $S_b = \{i\}$ for all $b \in B$.

If at least one competitive assignment is crowded we take it as one of P, P^*: this contradicts the existence of B and B^* and we are done. We continue the proof when we only know that preferences are semi-strict. Writing $T^* = \cup_{a \in B^*} S^*_a$ and $T = \cup_{a \in B^*} S_a$ we have $w_{T^*} = \sum_{a \in B^*} s^*_a > \sum_{a \in B^*} s_a = w_T$, implying that $T^* \diagdown T$ is not empty: it contains some agent $i \in S^*_a \cap S_b$ where $a \in B^*$ and $b \in C \cup B$. By definition of the demands and of the partition of A we have

$$(a, w_i) = (a, s^*_a) \succeq_i (b, s^*_b \vee w_i) \text{ and } (b, s_b) \succeq_i (a, s_a \vee w_i) = (a, w_i)$$

where the last equality is from $s_a < s^*_a = w_i$. If $b \in B$ we have $w_i = s_b > s^*_b$ therefore $(b, s^*_b \vee w_i) \simeq_i (b, s_b)$; this indifference still holds if $b \in C$ because in that case $s^*_b = s_b \geq w_i$: we conclude that all preferences above are indifferences in particular $(a, w_i) \simeq_i (b, s_b)$ which contradicts the semi-strictness of preferences.

Statement i) Unique competitive welfare. Fix $P = (S_x)_{x \in A}$ and $P^* = (S^*_x)_{x \in A}$ both competitive; we just proved $s = s^*$. An agent i is in some $S^*_a \cap S_b$ (where $a = b$ is possible); by definition of competitiveness and the inequalities $w_i \leq s_a, s_b$ we have $(a, s_a) \succeq_i (b, s_b)$ and $(b, s_b) \succeq_i (a, s_a)$.

Statement ii) Assume P is competitive and Pareto dominated by Q. Pick any agent i who is in S_a at P and in T_b at Q ($a = b$ is possible). We have

$$(b, t_b) \succeq_i (a, s_a) \succeq_i (b, t_b \vee w_i) \tag{20}$$

therefore $t_b \vee w_i \geq t_b$. Suppose $s_b < t_b$: taking into account $w_i \leq t_b$, the inequality $s_b \vee w_i \geq t_b$ implies $w_i = t_b > s_b$, and the two preferences in (20) are indifferences, in particular $(b, w_i) \simeq_i (a, s_a)$ a contradiction of semi-strict preferences if $a \neq b$. If $a = b$ then (20) gives $(b, t_b) \succeq_i (b, s_b)$ contradicting $s_b < t_b$.

We have shown $s_b \geq t_b$ everywhere on the support of Q, so these are all equalities. By replacing s_b by t_b in (20) we get $(b, t_b) \simeq_i (a, s_a)$ a final contradiction.

Statement iii) Fix a competitive assignment P, a post a in its support and an agent $i \in S_a$. We use the preferences $(a, s_a) \succeq_i (x, w_{S_x} \vee w_i)$ for all $x \neq a$ to evaluate the size of the set of allocations agent i may prefer to (a, s_a): at post a the size is $s_a - w_i$; at x such that $w_{S_x} \geq w_i$ the size is $w_{S_x} - w_i$; and 0 at x such that $w_{S_x} \leq w_i$. So the total is at most $W - w_i$. ■

In Example 6 (end of section 4.2) it is easy to check that P^1 is the unique competitive assignment. In the next Example there is no competitive assignment.

Example 7: $m = 2, n = 3$ *and* $W = 21$. The weights, $(W - w_i)$-prefixes, and the two top-$\frac{1}{2}$-fair assignments are

	a	b	w
α	16	15	10
β	16	15	10
γ	9	13	1

	a	b
P^1	α	$\beta\gamma$
P^2	β	$\alpha\gamma$

If the preferences are represented by the slack utilities as in Example 6 neither P^1 or P^2 is competitive because both α and β prefer $(a, 10)$ to $(b, 11)$.

Example 8: $m = 4, n = 4$. Every agent's priority is to minimise congestion: $s < s\prime$ implies $(x, s) \succ_i (y, s')$ for all posts x, y. The agents' preferences over two equally congested posts are as follows

$$\begin{array}{ll} \alpha & a \simeq b \simeq c \simeq d \\ \beta & b \succ a \succ c \succ d \\ \gamma & c \succ a \succ d \succ b \\ \delta & d \succ a \succ b \succ c \end{array}$$

so α's preferences are not semi-strict. Suppose $w_\beta = w_\gamma = w_\delta < w_\alpha$: this produces exactly four competitive assignments, identical in congestion but different in welfare

	a	b	c	d
P^1	α	β	γ	δ
P^2	β	α	γ	δ
P^3	γ	β	α	δ
P^4	δ	β	γ	α

Alternatively if α's preferences at constant congestion are also strict a single competitive assignment survives: the one most favourable to the heavier agent α, no matter how small is this advantage.

An example where both congestion and welfare profiles vary widely across two competitive assignments is the subject of section 6.3.

4.3 fractional competitiveness

Switching to cardinal vNM utility functions we partially emulate the results of section 3.2. Two of the key results are preserved: the existence of a competitive fractional congestion profile and its qualified uniqueness (Lemma 4); the implementation of the competitive congestion by one or more lotteries over deterministic assignments (Lemma 5). However we lose the rounding and approximation results of Lemma 3 and Theorem 1.

The congestion profile $\sigma = (\sigma_a)_{a\in A}$ is in the simplex $\Delta^{\mathbb{R}}(A, W)$ and each agent i has a cardinal vNM utility function $u_i(a, z)$ over $A \times [w_i, W]$, continuous and strictly increasing in z. As before if $\sigma_a < w_i$ the congestion price of post a to agent i is in fact w_i so we do not define $u_i(a, z)$ in the interval $0 \leq z < w_i$. Given a problem (A, N, u, w) agent i's competitive demand at σ is $D(u_i, \sigma) = \arg\max_{x\in A} u_i(x, \sigma_x \vee w_i)$ and the fractional congestion profile $\sigma \in \Delta^{\mathbb{R}}(A; W)$ is competitive (Comp)
iff $\sigma \in \sum_{i\in N} w_i \cdot \Delta^{\mathbb{R}}[D(u_i, \sigma); 1]$.

The (not necessarily unique) corresponding semi-stochastic matrix Π meets (16), but (15) is replaced by

$$\sigma_a = \sum_{i\in N} w_i \pi_{ia} \text{ for all } a \in A \tag{21}$$

The same fixed point argument as in Lemma 2 proves that a solution of (21) exists. To prove that it is unique we must strengthen the crowdedness property of Definition 7 and prove a weaker uniqueness statement than in Proposition 3.

Definition 8 *The fractional congestion profile σ is crowded in problem (A, N, u, w) iff for all a in its support ($\sigma_a > 0$) and each agent i demanding a ($a \in D(u_i, \sigma)$) we have $\sigma_a > w_i$.*

Lemma 4: *In any problem $(A, N, \succeq, w)$ there is at most one crowded competitive congestion profile.*

Proof Let σ,σ^* be two different crowded competitive congestion profiles, and B^* the non empty set of posts s. t. $\sigma^*_a > \sigma_a$. We claim that if $D(u_i, \sigma^*)$ contains a post $a \in B^*$ for some agent i then $D(u_i, \sigma) \subseteq B^*$. If, to the contrary, $D(u_i, \sigma)$ contains b outside B^* we get the familiar inequalities $u_i(a, \sigma^*_a) \geq u_i(b, \sigma^*_b \vee w_i)$ and $u_i(b, \sigma_b) \geq u_i(a, \sigma_a \vee w_i) > u_i(a, \sigma^*_a)$ where the strict inequality is from $\sigma_a < \sigma^*_a$ and the crowding assumption. If $b \notin B^*$ we have $\sigma^*_b \vee w_i \leq \sigma_b$ and derive a contradiction from $u_i(b, \sigma^*_b \vee w_i) \geq u_i(b, \sigma_b)$.

Pick now some semi-stochastic matrices Π and Π^* meeting (21) for σ and σ^* respectively, and use the property $D(u_i, \sigma^*) \cap B^* \neq \varnothing \Longrightarrow D(u_i, \sigma^*) \subseteq B^*$ to compute for any $i \in N$: $\sum_{a \in A^*} \pi^*_{ia} > 0 \Longrightarrow \sum_{a \in A^*} \pi^*_{ia} \leq \sum_{a \in A^*} \pi_{ia}$, which implies $\sum_{a \in A^*} \sigma^*_a = \sum_{i \in N} w_i(\sum_{a \in A^*} \pi^*_{ia}) \leq \sum_{i \in N} w_i(\sum_{a \in A^*} \pi^*_{ia}) = \sum_{a \in A^*} \sigma_a$, the final contradiction. ■

The reduced version of Lemma 3 is now a simple statement.

Lemma 5 *In any problem (A, N, u, w) we can implement any competitive congestion profile σ^c by one or more list $(\{P^k\}_{k=1}^K, \mathcal{L})$ of K deterministic assignments P^k together with a lottery $\mathcal{L}$ over $[K]$, such that their expected congestion is σ^c and they always assign the agents in their competitive demands: $i \in S^k_a \Longrightarrow a \in D(u_i, \sigma^c)$.*

Proof Pick Π realising the congestion σ^c in (21) and apply the following version of the Birkoff theorem: every semi-stochastic matrix Π is the convex combination of deterministic matrices (all entries are 0 or 1) of which the support is contained in that of Π. ■

In Example 7 (section 4.2) the fractional competitive congestion is $\sigma^c = (11, 10)$

because $D(\alpha, \sigma^c) = D(\beta, \sigma^c) = \{a, b\}$. It can be implemented in two different ways combining the top-$\frac{1}{2}$-fair assignments P^1, P^2 with some not top-$\frac{1}{2}$-fair assignments.

Specifically as $\mathcal{L}^1 = \frac{1}{4}P^1 + \frac{1}{4}P^2 + \frac{1}{4}P^3 + \frac{1}{4}P^4$, where

	a	b
P^3	$\alpha\gamma$	β
P^4	$\beta\gamma$	α

or as $\mathcal{L}^2 = \frac{9}{20}P^1 + \frac{9}{20}P^2 + \frac{1}{10}P^5$, where $P^5 = \begin{array}{cc} a & b \\ \alpha\beta & \gamma \end{array}$. The former keeps congestion close to top-$\frac{1}{2}$-fair, but violates this property half of the time; the latter only violates it with probability $\frac{1}{10}$ but then much more severely.

5 concluding comments

Take-home points.

Whether congestion is anonymous or weighted, the canonical guarantee eliminates for each participant all but the best $\frac{1}{m}$-th quantile of the feasible allocations, where m is the number of posts (Lemma 1 and Propositions 2).

Competitiveness, when it exists in the deterministic version of either model, identifies a single assignment in terms of congestion and welfare: combining efficiency with natural ex ante and ex post fairness properties, it is then a compelling normative solution to the congested assignment problem (Propositions 1 and 3).

The randomised competitive demand is unique (Lemmas 2 and 4) and implemented by a lottery over competitive deterministic assignments (Lemmas 3 and 5). In the anonymous case the latter are close to each other, as well as approximately top-n-fair, competitive, and efficient (Theorem 1).

Two open questions

1). The addition of post-specific lower and upper bounds on congestion is natural in many of the motivating examples discussed in the Introduction. Under anonymous congestion we can easily generalise the concept of top-n-fair guarantee but its interpretation is a bit more involved. For agent i the profile λ_i must satisfy (1) as well as the bounds: $\gamma_a^- \leq \lambda_{ia} \leq \gamma_a^+$ for all a. So the number of allocations that may or may not be accepted by λ_i is $\sum_a(\gamma_a^+ - \gamma_a^-)$ and λ_i captures the best $(n - \sum_a \gamma_a^-)$ of these. Then Lemma 1 goes through exactly as before.

If the lower bound γ_a^- forces the agents to populate a post a that they unanimously loath, our definition of competitiveness must be adapted to hope for an existence result. It is not clear how this can be done, whether in the deterministic or randomised model. Even with only upper bounds we lose the critical fixed point argument in Lemma 2 for the existence of a randomised competitive congestion profile.

2). Our approach can be applied to the "dual" domain of preferences where more congestion is strictly desirable: $(a, s+1) \succ (a, s)$ for all a and s. Interpret the report λ_i satisfying (1) for all a as "agent i accepts (a, s) only if $s \geq \lambda_{ia}$"[11]. Then Lemma 1 goes through: for any profile $(\lambda_i)_i$ of such reports there is an assignment satisfying these constraints. The proof mimicks that of Lemma 1 by switching a couple of signs.

The Definition 2 of a competitive assignment is unchanged but, even in the deterministic case, it no longer identifies a unique solution as explained in section 6.4. Proposing and justifying a fair refinement of the competetitive correspondence in the "good congestion" model is a challenging question, one we must answer before we can discuss the very natural domain of preferences single-peaked w. r. t. the congestion at each post.

References

[1] Abdulkadiroglu A. and T. Sonmez. 2003. School Choice: A Mechanism Design Approach, *American Economic Review*, 93 (3), 729–747.

[2] Banerjee S., Konishi, H. and Sonmez, T. 2001. Core in a Simple Coalition Formation Game, *Social Choice Welfare* 18, 135–153.

[3] Biro P., Fleiner T., Irving R.W. and Manlove, D. F. 2010. The college admissions problem with lower and common quotas, *Theoretical Computer Science*, 411, 3136–3153.

[11] This allows each agent to veto at most $n-1$ allocations, avoiding only the lowest $\frac{1}{m}$ quantile of his preferences.

[4] Bogomolnaia A. 2015. Random assignment: redefining the serial rule, *Journal of Economic Theory*, 158, 308-318.

[5] Bogomolnaia A. and M. O. Jackson. 2002. The stability of hedonic coalition structures, *Games & Economic Behavior,* 38,2, 201-230.

[6] Bogomolnaia A., Le Breton M., Savvateev A. and S. Weber. 2007. Stability under unanimous consent, free mobility and core, *International Journal of Game Theory*, 35,2,185-204.

[7] Bogomolnaia A. and H. Moulin. 2001. A new solution to the random assignment problem, *Journal of Economic Theory*, 100,2,201-230.

[8] Boiteux M. 1960. Peak-Load pricing, *The Journal of Business*, 33, 2,157-179.

[9] Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes, *J. Political Economy,* 119, 6, 1061-1103.

[10] Cacchiani V., Iori M., Locatelli M. and Martello S. 2022. Knapsack problems, an overview of recent advances, Part II: Multiple,multidimensional, and quadratic knapsack problems, *Computers and Operations Research*, 143, 105693.

[11] Caragiannis C., Kaklamanis C., Kanellopoulos P. and E. Papaioannou. 2008. Scheduling to maximize participation, *Theoretical Computer Science,* 402, 142-155.

[12] Copland A. 2021. School choice and class size externalities. http://dx.doi.org/10.2139/ssrn.4346860

[13] Correa J.R. and Stier-Moses N. E. 2010. Wardrop equilibria, *Wiley Encyclopedia of Operations Research and Management Science*, ed. James J. Cochran, John Wiley & Sons, Inc.

[14] Ergin H. and Sönmez T. 2006. Games of school choice under the Boston mechanism, *Journal of Public Economics*, 90, 215-237.

[15] Foley, D. K. 1966. *Resource allocation and the public sector*, Yale University.

[16] Frank M. 1981. The Braess paradox, *Mathematical Programming*, 20, 283-302.

[17] Green J. and J. J. Laffont. 1977. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods, *Econometrica*, 45, 2, 427-438.

[18] Hafalir I. E., Yenmez M. B. and M. A. Yildirim. 2013. Effective affirmative action in school choice, *Theoretical Economics*,8,2, 325-363.

[19] Hylland A. and R. Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions, *Journal of Political Economy*, 87, 2, 293-31.

[20] Kamada Y. and Kojima F. 2015. Efficient matching under distributional constraints: Theory and applications, *American Economic Review*, 105(1):67–99.

[21] Katta A-K. and J. Sethuraman. 2006. A solution to the random assignment problem on the full preference domain, *Journal of Economic Theory, 131, 1, 231-250.*

[22] Konishi H., Le Breton M., and S. Weber. 1997. Equilibrium in a Model with Partial Rivalry, *Journal of Economic Theory,* 72, 225]237.

[23] Koutsoupias E. and C. H. Papadimitriou. 1999.Worst-case equilibria, *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, 404-413.

[24] Milchtaich I.1996. Congestion games with player-specific payoff functions, *Games & Economic Behavior*, 13, 1, 111-124.

[25] Monderer D., and Shapley L. S. 1996. Potential Games, *Games & Economic Behavior,* 14, 1, 124-143.

[26] Monte D. and Tumennasan, N. 2013. Matching with quorums, *Economics Letters*, 120(1):14–17

[27] Moulin H. 2019. Fair Division in the Internet Age, *Annual Review of Economics,* 11, 407-41.

[28] Papai S. 2000. Strategyproof assignment by hierarchical exchange, *Econometrica,* 68, 6, 1403-1433.

[29] Phan W., Tierney R., and Y. Zhou. 2023. Crowding in School Choice, https://wphan.wordpress.ncsu.edu

[30] Pigou A.C. 1920. *The Economics of Welfare*, London, Macmillan.

[31] Ramsey, F. P. 1927. A Contribution to the Theory of Taxation, *Economic Journal*, 37, 47-61.

[32] Rosenthal R. W. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria, *Int. J. Game Theory* 2, 65–67.

[33] Roughgarden T. 2005. *Selfish Routing and the Price of Anarchy*, MIT Press, Cambridge Mass.

[34] Roughgarden T. and E. Tardos. 2002. How bad is selfish routing? *Journal of the ACM*, 49(2), 236-259.

[35] Shapley L. and H. Scarf. 1974. On Cores and Indivisibility, *Journal of Mathematical Economics*, 1, 23 37.

[36] Steinhaus H. 1948. The Problem of Fair Division. *Econometrica,* 16(1), 101-4.

[37] Varian H. 1974. Equity, Envy and Efficiency, *Journal of Economic Theory*, 29(2), 217-244.

[38] Vickrey W. S. 1969. Congestion Theory and Transport Investment, *American Economic Review (Papers and Proceedings)*, 59, 251–260.

6 appendix

6.1 FM equilibria with different congestion and welfare (section 3.1.1)

We have 5 posts and 26 agents. Agent i is single-minded on post a, and agent j on post e: $\lambda_{ia} = \lambda_{je} = 26$. The other four types have the following preferences:

$$\begin{array}{r l}
\text{four } \alpha\text{-s:} & (b,4) \succ_\beta (a,2) \succ_\beta (a,5) \succ_\beta (b,9) \\
\text{eight } \beta\text{-s:} & (c,8) \succ_\alpha (b,5) \succ_\alpha (b,8) \succ_\alpha (c,9) \\
\text{eight } \gamma\text{-s:} & (c,8) \succ_\gamma (d,5) \succ_\gamma (d,8) \succ_\gamma (c,9) \\
\text{four } \delta\text{-s:} & (d,4) \succ_\delta (e,2) \succ_\delta (e,5) \succ_\delta (d,9)
\end{array}$$

where if a post x is not listed for agent i it means that i's 26-prefix contains no allocation at this post. It is easy to check that the two following assignments are FM equilibrium outcomes:

	a	b	c	d	e
P^1	$i\alpha\alpha\alpha\alpha$	$\beta\beta\beta\beta\beta\beta\beta\beta$	$\gamma\gamma\gamma\gamma\gamma\gamma\gamma\gamma$	$\delta\delta\delta\delta$	j
P^2	i	$\alpha\alpha\alpha\alpha$	$\beta\beta\beta\beta\beta\beta\beta\beta$	$\gamma\gamma\gamma\gamma\gamma\gamma\gamma\gamma$	$j\delta\delta\delta\delta$

6.2 unbounded envy in the ordinal approach (section 3.2.2)

Consider Example 4 (section 3.2.1) where the two β-s have identical ordinal preferences that first minimise congestion then favour a over b:
$\cdots \prec_\beta (b,5) \prec_\beta (a,5) \prec_\beta (b,4) \prec_\beta (a,4) \prec_\beta (b,3) \prec_\beta (a,3) \prec_\beta \cdots$; also the four α-s are single-minded on post a and the two γ-s on post b. The unique top-8-fair assignment puts all α-s at a and the rest at b, therefore $s = (4,4)$ is the most natural choice of congestion.

In the corresponding decongested problem the two $\widetilde{\beta}$-s prefer $\widetilde{a}$ to $\widetilde{b}$ so they first share the four $\widetilde{a}$-s with the $\widetilde{\alpha}$-s before sharing with everyone what remains of $\widetilde{b}$. This produces the random assignment where each α and each β gets $\frac{2}{3}a + \frac{1}{3}b$, and the γ-s are always at b. Its implementation is $\frac{1}{3}P^1 + \frac{1}{3}R^1 + \frac{1}{3}R^2$ where

	a	b
P^1	$\alpha\alpha\alpha\alpha$	$\beta\beta\gamma\gamma$
R^1	$\alpha_1\alpha_2\beta\beta$	$\alpha_3\alpha_4\gamma\gamma$
R^2	$\alpha_3\alpha_4\beta\beta$	$\alpha_1\alpha_2\gamma\gamma$

We see that ex post at $R^{1,2}$ the unlucky α-s assigned to $(b, 4)$ are envious of the allocation $(a, 4)$ by at least seven welfare levels: single-mindedness implies
$(b, 4) \prec_\alpha (b, 3) \prec_\alpha \cdots \prec_\alpha (b, 1) \prec_\alpha (a, 8) \prec_\alpha \cdots \prec_\alpha (a, 4)$.

Now we replicate Example 4 by an integer factor of q where the α-s and γ-s are still single-minded at a and b respectively, while the β-s' preferences stil minimise congestion before favouring a over b:
$\cdots \prec_\beta (a, 4q + 1) \prec_\beta (b, 4q) \prec_\beta (a, 4q) \prec_\beta (b, 4q - 1) \prec_\beta \cdots$. The unique top-fair congestion is $s = (4q, 4q)$ and, as before, in the decongested problem the $\widetilde{\alpha}$-s and $\widetilde{\beta}$-s share the $\widetilde{a}$-s before they all finish the $\widetilde{b}$-s. The final random assignment still gives $\frac{2}{3}a + \frac{1}{3}b$ to the α-s and β-s, and b for sure to the γ-s; it is implemented by a fair lottery over the q-replication of P^1, R^1 and R^2. For instance in R^1 half of the α-s share $(a, 4q)$ with the β-s and the other half share $(b, 4q)$ with the β-s. Evaluating as before the level of envy for the unlucky α-s, we count $4q - 1$ welfare levels between $(b, 4q)$ and $(b, 1)$ and another $4q - 1$ between $(a, 8q)$ and $(a, 4q)$, for a total of $8q - 1$ or $8q - 2$.

This is in stark contrast with our competitive solution that for *any* profile of vNM utilities compatible with the ordinal preferences above finds a competitive congestion $\sigma^c = (4q + x, 4q - x)$ with $0 \leq x \leq 1$, because
the sign of $u_\beta(a, 4q + x) - u_\beta(b, 4q - x)$ changes between $x = 0$ and $x = 1$. Then, just as in section 3.2.1 for the original Example 4, the congestion σ^c is implemented by a lottery over the two assignments

	a	b
P^1	$4q \times \alpha$	$2q \times \beta, 2q \times \gamma$
P^2	$4q \times \alpha, \beta$	$(2q - 1) \times \beta, 2q \times \gamma$

and the β-s' preferences show that at P^2, the level of envy of the single unlucky β posted at a is just three; at P^1 it is just one for each β.

6.3 competitive assignments with different weighted congestion and welfare (section 4.2)

This example has $m = 5$ and 6 agents labeled α to η with the weights $w_\alpha = w_\beta = 4$; $w_\delta = w_\varepsilon = 3$; $w_\gamma = w_\eta = 1$. Their preferences are

$$\begin{gathered}\alpha\text{: } (b,4) \simeq (c,5)\\ \beta\text{: } (d,4) \simeq (c,5)\\ \gamma\text{: } (e,1) \succ (c,5) \succ (e,3)\\ \delta\text{: } (b,3) \succ (a,3) \succ (b,4)\\ \varepsilon\text{: } (d,3) \succ (e,3) \succ (d,4)\\ \eta\text{: } (a,1) \succ (c,5) \succ (a,3)\end{gathered}$$

where if a post x is not listed for agent i it means that $\lambda_{ix} = 0$. Only agents α and β's preferences are not semi-strict. It is straightforward to check that the following two assignments are competitive:

	a	b	c	d	e
P^1	δ	α	η, β	ε	γ
P^2	η	δ	α, γ	β	ε

congestion:

	a	b	c	d	e
P^1	3	4	5	3	1
P^2	1	3	5	4	3

The congestion of every post except c changes between the two assignments. Agents $1, 2$ are indifferent between P^1 and P^2; agents $3, 5$ strictly prefer P^1 and agents $4, 6$ strictly prefer P^2.

6.4 multiple competitive assignments under good congestion

Suppose $n \geq 2m$ and $\lambda_{ia} \geq 2$ for all i and a: then the m assignments where all agents share the same post are competitive.

For another simple example we have $m = 2$ and four agents with dual slack utilities $u_i(a, s) = s - \lambda_{ia}$. Their minimal acceptable congestions $(\lambda_i)_{i \in N}$ are

	a	b
$\alpha\alpha$	1	3
β	2	2
γ	3	1

meaning for instance that β only rejects $(a, 1)$ and $(b, 1)$ and γ only $(a, 1)$ and $(a, 2)$. There are six competitive assignments where post a gets successively $\alpha\alpha\beta\gamma$, $\alpha\alpha\beta$, $\alpha\alpha$, one α (two assignments), or post a is empty.